arXiv:1611.03745v1 [cs.DS] 11 Nov 2016

Dynamic Spanning Forest with Worst-Case Update Time: Adaptive, Las Vegas, and $O(n^{1/2-\epsilon})$-Time

Danupon Nanongkai
KTH Royal Institute of Technology

Thatchaphol Saranurak
KTH Royal Institute of Technology

Abstract

We present two algorithms for dynamically maintaining a spanning forest of a graph undergoing edge insertions and deletions. Our algorithms guarantee *worst-case update time* and work against an *adaptive adversary*, meaning that an edge update can depend on previous outputs of the algorithms. We provide the *first polynomial improvement* over the long-standing $O(\sqrt{n})$ bound of [Frederickson STOC'83, Eppstein, Galil, Italiano and Nissenzweig FOCS'92] for such type of algorithms. The previously best improvement was $O(\sqrt{n(\log\log n)^2/\log n})$ [Kejlberg-Rasmussen, Kopelowitz, Pettie and Thorup ESA'16]. We note however that these bounds were obtained by deterministic algorithms while our algorithms are randomized.

Our first algorithm is Monte Carlo and guarantees an $O(n^{0.4+o(1)})$ worst-case update time, where the $o(1)$ term hides the $O(\sqrt{\log\log n/\log n})$ factor. Our second algorithm is Las Vegas and guarantees an $O(n^{0.49305+o(1)})$ worst-case update time with high probability. Algorithms with better update time either needed to assume that the adversary is oblivious (e.g. [Kapron, King and Mountjoy SODA'13]) or can only guarantee an amortized update time. Our second result answers an open problem by Kapron et al. To the best of our knowledge, our algorithms are among a few non-trivial randomized dynamic algorithms that work against adaptive adversaries.

The key to our results is a decomposition of graphs into subgraphs that either have high expansion or sparse. This decomposition serves as an interface between recent developments on (static) flow computation and many old ideas in dynamic graph algorithms: On the one hand, we can combine previous dynamic graph techniques to get faster dynamic spanning forest algorithms if such decomposition is given. On the other hand, we can adapt flow-related techniques (e.g. those from [Khandekar, Rao and Vazirani STOC'06], [Peng SODA'16], and [Orecchia and Zhu SODA'14]) to maintain such decomposition. To the best of our knowledge, this is the first time these flow techniques are used in fully dynamic graph algorithms.

Contents

1 Introduction **1**

2 Overview **4**

3 Preliminaries **10**
3.1 Augmented ET Tree . . . 11
3.2 Definition of Dynamic SF . . . 12

4 Basic Tools **13**
4.1 2-dimensional ET Tree . . . 13
4.1.1 Reduction to a Constant Number of Connected Components . . . 14
4.1.2 List Intersection Oracle . . . 14
4.1.3 Non-tree Euler Tour . . . 15
4.2 Cut Recovery Tree . . . 16
4.2.1 Sparse Recovery . . . 17
4.2.2 d-word ET Tree . . . 17
4.2.3 Put Them Together . . . 18
4.3 Classic Reductions . . . 19
4.3.1 Overview of the Sequence of Reductions . . . 20

5 Global Expansion Decomposition **23**
5.1 Most-Balanced Sparse Cut . . . 24
5.2 The Global Decomposition Algorithm . . . 24

6 Local Expansion Decomposition **28**
6.1 Most-balanced Sparse Cut with respect to Overlapping Cuts . . . 29
6.2 The Local Decomposition Algorithm . . . 30
6.2.1 Validity . . . 30
6.2.2 Correctness . . . 31
6.2.3 Running time . . . 34

7 Monte Carlo Dynamic SF with $O(n^{0.4+o(1)})$ Update Time **35**
7.1 Dynamic SF on Expanders . . . 36
7.2 Reduction to Expanders via Global Expansion Decomposition . . . 37
7.2.1 Reduction for Las Vegas algorithms . . . 39

8 Las Vegas Dynamic SF with $O(n^{0.49305+o(1)})$ Update Time **40**
8.1 Dynamic SF on Broken Expanders via Local Expansion Decomposition . . . 40
8.1.1 Preprocessing . . . 41
8.1.2 Handling Updates . . . 44
8.2 Reduction to Broken Expanders . . . 46

A Formalization of Dynamic Algorithms vs. Adversaries **52**

B Omitted proof **54**
B.1 Proof of Theorem 3.3 . . . 54
B.2 From Section 4.1 and Section 4.2 . . . 54
B.2.1 Proof of Theorem 4.5 . . . 54
B.2.2 Proof of Theorem 4.14 . . . 56
B.2.3 Proof of Theorem 4.16 . . . 57
B.3 From Section 4.3 . . . 58
B.3.1 Proofs of Lemma 4.23, Lemma 4.24, Lemma 4.25 and Lemma 4.26. . . . 58
B.3.2 Proofs of Lemma 4.27 and Lemma 4.28 . . . 61
B.4 Proof of Theorem 5.4 . . . 64
B.5 Proof of Theorem 6.5 . . . 67

C Very Simple Dynamic SF with $O(m^{0.5+o(1)})$ Update Time **70**
C.1 The Algorithm . . . 70
C.2 Analysis . . . 70

1 Introduction

In the *dynamic spanning forest (*SF*)* problem, we want to maintain a spanning forest F of an undirected unweighted graph G undergoing edge insertions and deletions. In particular, we want to construct a data structure that supports the following operations.

- PREPROCESS(G): Initialize the data structure with an input graph G. After this operation, the data structure outputs a spanning forest F of G.
- INSERT(u, v): Insert edge (u, v) to G. After this operation, the data structure outputs changes to F, if any.
- DELETE(u, v): Delete edge (u, v) from G. After this operation, the data structure outputs changes to F, if any.

The goal is to minimize the *update time*, i.e. the time needed to execute the insert and delete operations. The time bound is in terms of n and m, denoting respectively the number of nodes and edges in G at the time the update happens. We use $\tilde{O}$ to hide $\operatorname{poly}\log(n)$ terms.

The update time is usually categorized into two types: For any t, an algorithm is said to have an *amortized* update time of t if, for any k, the total time it spends to process the first k updates (edge insertions/deletions) is at most kt. It is said to have a *worst-case* update time of t if it spends at most t time for *each* update operation. (In case of randomized algorithms, the guarantee may hold with high probability or in expectation.) Thus, roughly speaking an algorithm with a small amortized update time is fast "on average" but may take a long time to respond to a single update. In contrast, the worst-case update time is more preferable since it guarantees to hold for every operation. (We also note that the amortized update time is less desired because most algorithms that guarantee it have to assume that the input graph has no edge initially.)

The dynamic SF problem along with its closely related variations – *dynamic connectivity* and *dynamic minimum spanning forest* (MSF) – played a central role in the study of dynamic graph algorithms. The first result for these problems dates back to Frederickson's deterministic algorithm from 1985 [12], which provides a worst-case $O(\sqrt{m})$ update time. This bound, when combined when the general sparsification technique of Eppstein et al. [11] from 1992, implies a worst-case $O(\sqrt{n})$ update time. By allowing the update time to be amortized, this bound was significantly improved by Henzinger and King [16] in 1995, who presented a Las Vegas randomized algorithm with polylogarithmic expected amortized update time. In the following decade, this result was refined in many ways, including algorithms with smaller update time which almost matches existing lower bounds (the gap is currently $O((\log\log n)^2)$; see, e.g., [20, 38, 18, 31]), deterministic algorithms (e.g. [19, 41]), and algorithms that works for harder problems such as two-edge connectivity and MSF (e.g. [19, 17]).

Given that the problem is fairly well-understood from the perspective of amortized update time, many researchers have turned their attention back to the worst-case update time in the last decade (one sign of this trend is the 2007 work of Pătraşcu and Thorup [32]). The question on worst-case update time is not only interesting in the context of the dynamic SF problem. This is because for many problems (e.g. MSF, two-edge connectivity, shortest paths [39] and matching [8]) we have many techniques to argue about the amortized update time but do not understand the case of worst-case update time much. In the context of dynamic SF, the $O(\sqrt{n})$ worst-case update time of [12, 11] has remained the best for decades until Kapron, King and Mountjoy [23] showed a breakthrough polylogarithmic bound in 2013 (the bound was originally $O(\log^5 n)$ in [23] and was later improved to $O(\log^4 n)$ [13]).

We however still cannot claim a victory over this problem, because the algorithms of [23, 13] still need to be refined in many ways. Besides the possibility to improve the bound further, one big challenge arose from the fact that they are randomized *Monte Carlo*, meaning that they may make mistakes (with a small probability). While the ultimate goal would be to get a deterministic algorithm with a similar worst-case update time that works for harder problems (such as two-edge connectivity and MSF), the next challenges are to obtain (i) a *Las Vegas* algorithm (thus making no errors) and (ii) an algorithm that works against *adaptive adversaries* (any deterministic algorithm will guarantee both of these). Getting *Las Vegas* algorithms is a well-known and interesting question in many areas of theoretical computer science. Since the issue about adaptive adversaries might be quite specific to the area of dynamic algorithms, we discuss this issue briefly here.

An adaptive adversary is the one who can determine an update based on previous outputs of the algorithm (e.g. the maintained spanning forests). It is used to contrast the weaker notion of *oblivious adversaries* who must fixed edge updates in advance and thus updates are not influenced by algorithms' outputs; for example, dynamic SF algorithms that work under the oblivious adversary assumption cannot reveal the maintained forest to an adversary. (Note that both types of adversaries do not have an access to the randomness used by the algorithm to make random choices[1].) One drawback of the algorithms of [23, 13] is that they work only under this assumption. This drawback in fact reflects our lack of understanding in exploiting randomness for a wide range of dynamic graph problems: without this assumption *very few* algorithms are able take advantage of randomness[2] It is a fundamental question whether the *true source of power* of randomized dynamic algorithms is the randomness itself or in fact the oblivious adversary assumption. We note that this question is also *very important in some applications*. For example, removing the oblivious adversary assumption from the randomized algorithms of Henzinger et al. [14] and Roditty-Zwick [34] for the decremental single-source and all-pairs shortest-paths problems will lead to new fast static algorithms for computing maximum s-t and multicommodity flows (see, e.g., [6, 28] for further discussions). This is the main drive behind the current efforts in *derandomizing* the algorithms of [14, 34] (e.g., [6, 7, 15]). (Note that deterministic algorithms always work against adaptive adversaries.) Another example is the reduction by Henzinger and King [16] to obtain a dynamic *approximate* minimum spanning forest algorithm $\mathcal{A}$ from a dynamic spanning forest algorithm $\mathcal{B}$. This reduction requires that $\mathcal{B}$ works against adaptive adversaries, even if we just want $\mathcal{A}$ to only work against oblivious adversaries[3]. Thus, we cannot directly apply this reduction to the algorithms of [23, 13] to obtain a dynamic approximate minimum spanning forest algorithm with polylogarithmic worst-case update time[4].

[1] In fact, one can categorize adversaries further based on how they can access the randomness. For example, we can define an *randomness-oblivious* adversary to be the one that never see the random bits and an *randomness-adaptive* adversary to be the one that see the random bits *after* they are used by the algorithm. We are not aware of algorithms that need the distinction between the two cases, so we do not discuss them here. We note that our algorithms work against the stronger notion of randomness-adaptive adversaries. Also note that one can define an even stronger notion of adversary who can see the entire random string from the beginning. Similar to the case of online algorithms [4], it can be argue that randomization does not help against this type of adversary; i.e. if there is a randomized algortihm that works against such adversary, then there is also a deterministic algorithm with the same performance.

[2] To the best of our knowledge, the only randomized dynamic graph algorithm that works against an adaptive adversary is that by [16] and the follow-up improvement by [38] and [24] which are also against adaptive adversary for the same reason. This excludes obvious cases where problem outputs are unique and thus there is no difference between adaptive and oblivious adversary for the problem.

[3] In particular, the reduction of Henzinger and King [16] involves maintaining a spanning forest F_i on graph G_i consisting of edges of weight roughly $(1+\epsilon)^i$ in the input graph and the maintained minimum spanning forest F, for every i. Consequently, F_i may affect F, which in turns affect the change G_i. So, the algorithm that maintains F_i must be able to handle an adaptive adversary since its own output (F_i) may influence its input (G_i).

[4] Note that Kapron et al. [23] and Gibb et al. [13] claimed that the reduction of Henzinger and King [16] together with their algorithms imply dynamic algorithms for the approximate minimum spanning forest problem with polylogarithmic worst-case update time. We believe that this claim is not correct. However, we also believe that one can

Motivated by the discussions above, we focus on scenarios where algorithms can guarantee worst-case update time and additionally either (i) work against adaptive adversaries, or (ii) are Las Vegas, or both. To this end, we note again that deterministic algorithms can guarantee both (i) and (ii). Thus, the classic $O(\sqrt{n})$ bound of [11, 12] applies to this scenario. This bound was slightly improved to $O(\sqrt{n(\log\log n)^2/\log n})$ recently by Kejlberg-Rasmussen et al. [24] using word operations. It remains an important open problem whether we can *polynomially* improved the long-standing $O(\sqrt{n})$ bound as this will likely need new techniques beyond word tricks.

Our Contributions. We show two algorithms that achieve the above. Our first algorithm is a randomized *Monte Carlo* algorithm that works against adaptive adversaries and has $O(n^{0.4+o(1)})$ worst-case update time, where the $o(1)$ term hides the $O(\sqrt{\log\log n/\log n})$ factor. At any point in time the forest maintained by it is a spanning forest with high probability.

Theorem 1.1. *There is a randomized dynamic* SF *algorithm for any graphs with n nodes and m initial edges that works correctly against adaptive adversaries with high probability and has preprocessing time $O(m^{1+o(1)})$ and worst-case update time $O(n^{0.4+o(1)})$.*

The term "works correctly against adaptive adversaries with high probability" is a technical term which is formalized in Section 3 and Appendix A. Roughly it means that if we pick an arbitrary time t, and consider the solution maintained by the algorithm at that time, such solution is a spanning forest with high probability (when considered all possible random choices). Note that we are talking about an infinite number of updates here, so the algorithm might make several mistakes in the past (which the adversary can see). Regardless, the guarantee holds with high probability.[5]

Our second algorithm is a randomized *Las Vegas* algorithm that has $O(n^{0.49305+o(1)})$ worst-case update time with high probability (this also implies the same expected time). This algorithm also works against adaptive adversaries.

Theorem 1.2. *There is a randomized dynamic* SF *algorithm for any graphs with n nodes and m initial edges that works correctly against adaptive adversaries with certainty and has preprocessing time $O(m^{1+o(1)})$ and worst-case update time $O(n^{0.49305+o(1)})$ with high probability.*

Note that both algorithms are among a few randomized algorithms that work against adaptive adversaries. Moreover, the second result answer the open problem raised by Kapron et al. [23].

The key to our results is the notion of *expansion decomposition* which is a decomposition of a graph into a sparse graph and connected components with high expansion. This decomposition serves as an interface between recent flow-related techniques and known dynamic graph algorithmic techniques: On the one hand, we show how to efficiently construct and maintain the decomposition using flow-related techniques such as fast (static) max-flow approximation [33, 35, 25], cut-matching games [26], and local cut improvement [30]. On the other hand, the decomposition allows us to focus on solving the dynamic SF problem only on a very sparse graph and a graph with high expansion. In these cases, we can combine known techniques in dynamic graph algorithms such as sparse recovery[6] [13, 23, 5], sparsification [11], ET tree [16], a modification of the reduction from k-weight MSF to SF [16], and a variant of the 2-dimensional topology tree [12, 40]. We refer to Section 2 for a more comprehensive overview.

still obtain the latter result both by modifying the reduction of [16] and by modifying the algorithms of [23, 13].

[5] An intuition is that mistakes from the far past should not affect an algorithm since it has enough time to run a static algorithm to detect and fix such mistakes. Thus only the algorithm' outputs from a polynomial number of previous steps can affect the algorithm, but then we can use the union bound to argue that the algorithms did not make mistakes in any of these steps.

[6] Techniques for the 1-sparse recovery problem was used in [13, 23]. In this paper, we need more general techniques for the s-sparse recovery problem for large s.

To the best of our knowledge, our results are the first applications of flow-related techniques in fully-dynamic graph algorithms. (Previously Pătraşcu and Thorup [32] used some relevant techniques in a special setting where updates can happen (in bulk) only once.) These results suggest the possibility to obtain stronger results (e.g. lower update time, deterministic algorithms and algorithms for harder problems) by further understanding these techniques in the context of dynamic graphs.

2 Overview

To simplify our discussion, in this section we use the following notations. For any functions $f(n)$ and $g(n)$ of n, we say that $f(n) = \hat{o}(g(n))$ if there exists some constant $\epsilon > 0$ such that $f(n) = O(g(n)^{1-\epsilon})$. In this section we will focus on getting an $\hat{o}(n^{1/2})$ worst case update time; in other words, we focus on getting an $O(n^{1/2-\epsilon})$ worst-case update time for some very small constant $\epsilon > 0$ without worrying about the specific value of ϵ. First, it can be shown, using standard techniques, that we only have to focus on a special case as in the following lemma.

Lemma 2.1 (Details in Section 4.3)**.** *It is sufficient to construct a dynamic spanning forest algorithm that (i) has $n^{1+o(1)}$ proprocessing time, (ii) can handle a dynamic graph G that has maximum degree $\Delta \leq 3$ at all time, and (iii) can handle $\tau = n^{1/2+\epsilon}$ edge updates for some small constant $\epsilon > 0$ (as oppose to handling infinite number of updates).*

Proof Idea. First, we can assume that the graph is sparse by the sparsification technique of Eppstein et al. [11][7]. Next, we reduce to the case where the graph has maximum degree 3, which is also done implicitly in [12, 24]. Given a graph G, we maintain a graph G' with maximum degree 3 by "splitting" each node in G into a path in G' of length equals to its degree. As G is sparse, there are $O(n)$ nodes in G'. We assign weight 0 to edges in those paths, and 1 to original edges. Observe that tree-edges with weight 1 in a 2-weight minimum spanning forest F' of G' form a spanning forest F of G. So we are done by maintaining F'. (There is a reduction from dynamic k-weight MSF to dynamic SF by Henzinger and King [16].)

Next, we reduce to the case when there is only $\tau = n^{1/2+\epsilon}$ updates. We divide a long sequence of updates into phases of length τ. We concurrently maintain two instances $\mathcal{A}_1$ and $\mathcal{A}_2$ of a dynamic SF algorithm. In odd phases, $\mathcal{A}_1$ maintains a spanning forest with update time $O(n^{1/2-\epsilon})$, and in the same time we evenly distribute the work for preprocessing $\mathcal{A}_2$ into each update of this phase, which takes $O(n^{1+o(1)}/\tau)$ time per update. So, in the next even phase, $\mathcal{A}_2$ is ready to maintain a spanning forest. Then we do everything symmetrically in even phases. Finally, we can "combine" the two dynamic spanning forests maintained by $\mathcal{A}_1$ and $\mathcal{A}_2$ into one dynamic spanning forest using additional $O(\log^2 n)$ update time via a standard trick (which is implicit in a reduction from dynamic k-weight MSF to dynamic SF by [16]). In total, the update time is $O(n^{1/2-\epsilon} + n^{1+o(1)}/\tau + \log^2 n) = O(n^{1/2-\epsilon+o(1)})$[8]. □

Our technical ideas evolve around the notion of *graph expansion*. Given a graph $G = (V, E)$ and a set of nodes $S \subset V$, let $\partial_G(S)$ denote the set of edges across the cut S in G. The expansion of S

[7] One thing we have to be careful of is that the sparsification technique of [11] only works if we can maintain a *stable* spanning forest. The spanning forests maintained by our algorithms are *not* stable as exactly defined in [11]. This is mainly because our algorithms are randomized. However, this can be easily dealt with by slightly modifying the original notion of stability, as done in Section 4.3.

[8] Using this technique for randomized algorithms, we can only reduce from the case of the polynomial-length update sequences. This is because the trick for "combining" two dynamic spanning forests will accumulate the failure probability over time. To handle infinite-length sequences, we need a different reduction. See 4.28.

is $\phi_G(S) = |\partial_G(S)|/\min\{|S|, |V \setminus S|\}$, and the expansion of a graph G is $\phi(G) = \min_{\emptyset \neq S \subset V} \phi_G(S)$.[9] Each of our algorithms consists of two main components. The first component is algorithms to *decompose* a graph into subgraphs with high expansion together with some other subgraphs that are easy to handle. This allows us to focus only on high-expansion subgraphs. Maintaining a spanning forest in a high-expansion is our second component. To be concrete, let us start with our first decomposition algorithm which is a building block for both our results. This algorithm takes near-linear time to decompose an input graph into a sparse subgraph and many connected components with high expansion, where the sparsity and expansion are controlled by a given parameter α as in the following theorem.

Theorem 2.2. *There is a randomized algorithm $\mathcal{A}$ that takes as inputs an undirected graph $G = (V, E)$ with $n \geq 2$ vertices and m edges and a parameter $\alpha > 0$ (α might depend on n). Then, in $O(m^{1+o(1)})$ time, $\mathcal{A}$ outputs two graphs $G^s = (V, E^s)$ and $G^d = (V, E^d)$ with the following properties.*

- *$\{E^s, E^d\}$ is a partition of E,*
- *G^s is sparse, i.e. $|E^s| \leq \alpha n^{1+o(1)}$, and*
- *with high probability, each connected component C of G^d either is a singleton or has high expansion, i.e. $\phi(C) \geq \alpha$.*

Proof Idea (details in Section 5). There are two main steps. The first step is to devise a near-linear time approximation algorithm for a problem called *most balanced sparse cut*. This problem is to find a cut S with largest number of nodes such that $|S| \leq n/2$ and $\phi(S) \geq \alpha$. This problem is closely related to sparsest cut and balanced cut problems (cf. [27]). One way to approximately solve both problems is by using the *cut-matching game* framework of Khandekar, Rao and Vazirani [26] together with known exact algorithms for the maximum flow problem. We modify this framework (in a rather straightforward way) so that (i) it gives a solution to the most balanced sparse cut problem and (ii) we can use *approximate* maximum flow algorithms instead of the exact ones. By plugging in near-linear time max flow algorithms [33, 35, 25], our algorithm runs in near-linear time.

The second step is to use the most balanced sparse cut algorithm to construction the decomposition. We note that a similar decomposition was constructed by Andoni et al. [3] in the context of graph sketch. This algorithm repeatedly finds and removes a cut S such that $\phi(S) < \alpha$ from the input graph. It is too slow for our purpose since it may involve many cuts S such that $|S| = 1$, causing as many as $\Omega(n)$ repetitions and as much as $\Omega(mn)$ running time. Instead, we use the most balanced cut S. The idea is that if S is large, the cut divides the graph into two subgraphs with similar number of nodes, and this cannot happen more than $O(\log n)$ time. It is still possible that we find a cut S that $|S|$ is small, but we can argue that this does not happen often. To this end, we note that the exact algorithm and analysis is quite involved because we only have approximate guarantees about the cut. This incurs a factor of $n^{O(\sqrt{\log \log n / \log n})} = n^{o(1)}$ in the running time. □

We call the above algorithm the *global expansion decomposition algorithm* to contrast it with another algorithm that is *local* in the sense that it does not read the whole graph (this algorithm will be discussed soon).

[9] Readers who are familiar with the notions of expansion and conductance may observe that the two notions can be used almost interchangeably in our case since we can assume that our input graph has maximum degree at most three.

The Monte Carlo Algorithm (Details in Section 7). With the above decomposition, we already have one main subroutine for our Monte Carlo algorithm against adaptive adversaries. This subroutine is run at the preprocess with parameter $\alpha = 1/n^{\epsilon}$, for some constant $\epsilon > 0$, to decompose the input graph into G^d and G^s. Now we argue that the only other main subroutine that we will need is maintaining a spanning forest on G^d; in other words, we do not have to worry so much about the graph G^s. The intuition is that G^s has $O(\alpha n^{1+o(1)}) = \hat{o}(n)$ edges, and thus we can maintain a spanning tree on G^s in $\hat{o}(n^{1/2})$ worst-case update time using Frederickson's algorithm. (The real situation is slightly more complicated since we have to maintain a spanning tree in a graph G' consisting of edges in G^s *and* a spanning forest of G^d. This is because we need to get a spanning forest of $G = (V, E^d \cup E^s)$ in the end. Fortunately, although G' may not have $\hat{o}(n)$ edges, we only need to slightly modify Frederickson's data structure to obtain the $\hat{o}(n^{1/2})$ update time as desired.) Similarly, since we can assume that there are at most $\tau = \hat{o}(n)$ edges inserted due to Lemma 2.1, it will not be hard to handle edge insertions.

So, we are left with maintaining a spanning forest on each connected component C of G^d when it undergoes (at most τ) edge deletions. First we need another tool that is a simple extension of the tool used by Kapron et al. [23, 13]. In [23, 13], Kapron et al. showed a deterministic data structure that can maintain a forest F (not necessarily spanning) in a dynamic graph G and answer the following query: Given a pointer to one tree T in F, if $\partial_G(V(T))$ contains *exactly* one edge, then output that edge (otherwise, the algorithm can output anything). The update and query time of this data structure is polylogarithmic. The main tool behind it is a 1-*sparse recovery* algorithm from data streams (e.g. [5]). By using an s-sparse recovery algorithm instead for a parameter s, this data structure can be easily extended to answer the following query: Given a pointer to one tree T in F, if $\partial_G(V(T))$ contains *at most s edges*, then output all those edges. The update and query time of this data structure is $\tilde{O}(s)$. We call this data structure a *cut recovery tree*. (See Section 4.2 for details.) Now the algorithm:

Algorithm $\mathcal{B}$. Recall that initially $\phi(C) \geq \alpha$, and we want to handle at most τ edge deletions. Consider a spanning forest F of C maintained at any point in time. Consider when an edge e is deleted. If e is not in F, we do nothing. If it is, let T be the tree that contains e. The deletion of e divides T into two trees, denoted by T_1 and T_2. Our job is to find one edge in $\partial_C(V(T_1))$, if exists, to reconnect T_1 and T_2 in F. Assume wlog that $|V(T_1)| \leq |V(T_2)|$. First we sample $(1/\alpha + \sqrt{\tau/\alpha})\text{polylog}(n)$ edges among edges in C incident to nodes in T_1. This can be done in $\tilde{O}(1/\alpha + \sqrt{\tau/\alpha})$ worst-case time by ET tree just as Henzinger and King did in [16]. If one of the sampled edges connects between T_1 and T_2, then we are done. If not, we use the cut recovery tree with parameter $s = \sqrt{\tau/\alpha}$ to list (at most s) edges in $\partial_C(V(T_1))$. If some edge in $\partial_C(V(T_1))$ is listed, then we use such edge to reconnect T_1 and T_2; otherwise, we leave T_1 and T_2 as two separate trees in F.

The above algorithm takes worst-case update time $\tilde{O}(1/\alpha + \sqrt{\tau/\alpha}) = \tilde{O}(n^{1/4+O(\epsilon)})$, which is $\hat{o}(n^{1/2})$ for small enough ϵ. For the correctness analysis, we argue that with high probability the algorithm does not make any mistake at all over the period of τ updates. First note that we initially get the decomposition as promised by Theorem 2.2 with high probability. So we will assume that this is the case, and in particular $\phi(C) \geq \alpha$ initially. Now consider three cases.

- Case 1: $|V(T_1)| \geq 2\tau/\alpha$. In this case, at the time we delete e we know that

$$\begin{aligned} |\partial(V(T_1))| &\geq \alpha|V(T_1)| - \tau && \text{(since we have deleted at most } \tau \text{ edges in total)} \\ &\geq \alpha|V(T_1)|/2 && \text{(since } |V(T_1)| \geq 2\tau/\alpha). \end{aligned} \tag{1}$$

  Since there are at most $3|V(T_1)|$ edges incident to nodes in T_1, when we sample $(1/\alpha +$

$\sqrt{\tau/\alpha}$)polylog(n) edges in the first step, we will get one of the edges in $\partial(V(T_1))$ with high probability.

- Case 2: $|V(T_1)| < 2\tau/\alpha$ and $|\partial(V(T_1))| \geq \sqrt{\tau/\alpha}$ at the time we delete e. Similarly to the previous case, since there are at most $3|V(T_1)| < 6\tau/\alpha$ edges incident to nodes in T_1, when we sample $(1/\alpha + \sqrt{\tau/\alpha})$polylog($n$) edges in the first step, we will get one of the edges in $\partial(V(T_1))$ with high probability.
- Case 3: $|V(T_1)| < 2\tau/\alpha$ and $|\partial(V(T_1))| < \sqrt{\tau/\alpha}$ at the time we delete e. Since $|\partial(V(T_1))| < \sqrt{\tau/\alpha}$ the cut recovery tree data structure with parameter $s = \sqrt{\tau/\alpha}$ will list all edges in $\partial(V(T_1))$ correctly in the second step of the algorithm. Thus we will always find an edge to reconnect T_1 and T_2, if there is one.

To conclude, in all cases the algorithm will find an edge to reconnect T_1 and T_2, if there is one, with high probability. Observe that the analysis exploits the fact that C has high expansion (compared to τ) initially by arguing (as in Case 1) that if $V(T_1)$ is large, then there will be plenty of edges in $\partial(V(T_1))$ that are not deleted.

Finally, we provide some intuition why the analysis above holds even when we allow an adversary to see the maintained spanning forest (the adaptive adversary case). In fact, we argue that this is the case even when the algorithm *reveals all random choices* it has made so far to the adversary. In particular, we allow the adversary to see (i) all $(1/\alpha+\sqrt{\tau/\alpha})$polylog($n$) edges sampled in Step 1 and (ii) the initial decomposition (from the algorithm in Theorem 2.2). We can reveal (i) because every random bit is used *once* to sample edges, and will not be used again in the future. Thus, knowing these random bits is useless for the adversary to predict the algorithm's behavior in the future. We can reveal (ii) because our three-case analysis only needs to assume that the decomposition works correctly initially. The only thing the adversary can exploit is when this is not the case, but this happens with a small probability. We refer to Section 7 for formal arguments and further details.

Remark: We note a lesson that might be useful in designing randomized dynamic algorithms against adaptive adversaries in the future. The reason that most randomized algorithms fail against adaptive adversaries is that their future behavior heavily depends on random bits they generated in the past. In contrast, the random bits our algorithm used for edge sampling are not reused, thus do not affect the future. The dynamic SF algorithm by Henzinger and King [16] works against adaptive adversaries by exactly doing this. In addition to this, in our case, the random bits our decomposition algorithm used may affect the future since the decomposition is used throughout. However, what matters in the future is only whether the decomposition algorithm gives a correct output or not. How the output looks like does not really matter as long as it is correct, which is the case with high probability.

The Las Vegas Algorithm (Details in Section 8). The goal now is to construct an algorithm that can *detect when it makes mistakes* (so that it can, e.g., restart the process). To motivate our new algorithm, let us re-examine the previous Monte Carlo algorithm when it makes mistakes. First, the initial decomposition might not be as guaranteed in Theorem 2.2, and it is not clear how we can check this. We will leave this issue aside for the moment (it will be easy to handle once we take care of other issues). Now assuming that the decomposition is correct, another issue is that Algorithm $\mathcal{B}$ may also make errors as it can happen that none of the $(1/\alpha + \sqrt{\tau/\alpha})$polylog($n$) sampled edges are in $\partial(V(T_1))$, but there is actually an edge in $\partial(V(T_1))$ to reconnect T_1 and T_2. However, Case 1 of the analysis, which is the crucial part that exploits the fact that C has high expansion (compared to τ) initially, is still useful: in this case, we know that there *is* an edge to reconnect T_1 and T_2 since $|\partial(V(T_1))| \geq \alpha|V(T_1)|/2$ (Equation (1)). Thus in this case, the algorithm *knows* that it makes a mistakes if it does not find an edge to reconnect among the sampled edges.

When $|V(T_1)| < 2\tau/\alpha$ however, we do not know how to distinguish between Cases 2 and 3, and thus will naively consider all edges incident to nodes in C. Using this idea with parameters slightly adjusted, we have the following algorithm.

Algorithm $\mathcal{C}$. *Case 1:* $|V(T_1)| \geq 2\tau/\alpha$. Sample $\text{polylog}(n)/\alpha$ edges from edges incident to nodes in T_1. If one of these edges are in $\partial(V(T_1))$, then we can reconnect T_1 and T_2; otherwise, the algorithm realizes that it fails and outputs "fail". *Case 2:* Consider all edges incident to nodes in T_1. (Note that there are at most $3|V(T_1)| = O(\tau/\alpha)$ such edges.) If one of these edges are in $\partial(V(T_1))$, then we can reconnect T_1 and T_2; otherwise, T_1 and T_2 become separate trees in the maintained forest.

The above algorithm is clearly Las Vegas, as it realizes when it makes mistakes. This is also the case even when we take into account the fact that the initial decomposition may fail: when Algorithm $\mathcal{C}$ outputs "fail" it means that either (i) $\mathcal{C}$ does not get an edge in $\partial(V(T_1))$, which is guaranteed to exist if the decomposition is correct, as a sample, or (ii) the decomposition itself is incorrect that such edge does not exists. In other words, the output "fail" of $\mathcal{C}$ capture the failures of its own and of the decomposition algorithm.

Now, observe that the update time of $\mathcal{C}$ is dominated by the second step, which is $\tilde{O}(\tau/\alpha)$. This is *not* the $\hat{o}(n)$ as we desire (recall that $\tau = n^{1/2+\epsilon}$ and $\alpha = 1/n^{\epsilon}$). Unfortunately, we do not know how to improve $\mathcal{C}$'s update time. So, instead we have to be more clever in using $\mathcal{C}$. The plan is the following. We will divide the sequence of at most τ updates on each high-expansion component C into *phases*, where each phase consists of $\tau' = \hat{o}(n^{1/2})$ updates. In the beginning of each phase, we decompose C further into components whose expansion is $\alpha' \leq \alpha$. (Ideally, we want α' to be as high as α, but there are some limits to this, as we will see shortly.) In each component, say C', we run algorithm $\mathcal{C}$ as before, but with parameters α' and τ'. The correctness of the algorithm is guaranteed in the same way as before because C' has expansion at least α' initially. The worst-case update time is then $\tilde{O}(\tau'/\alpha')$. Now, observe that if the decomposition algorithm run in the beginning of each phase take λ time, then this cost can be charged to τ' updates in the phase, causing the cost of λ/τ' on average. By a standard technique (similar to the proof sketch of Lemma 2.1), we can turn this averaged time into a worst-case update time. Thus, the total worst-case update time of the algorithm will be

$$\tilde{O}(\tau'/\alpha' + \lambda/\tau'). \tag{2}$$

Our goal is to make the above update time be $\hat{o}(n^{1/2})$. To do this, there is one obstruction though: If we use the decomposition algorithm from Theorem 2.3 in the beginning of every phase, then λ can be as large as $O(n^{1+o(1)})$ (when the connected component is big). It will then be impossible to make the above update time be $\hat{o}(n^{1/2})$ (note that $\tau' = \hat{o}(n^{1/2})$). So, we need a different algorithm for the initial decomposition of each phase, and in particular it should not read the whole connected component. The last piece of our algorithm is such a decomposition algorithm, which we call a *local algorithm* as it does not need to read the whole connected component we are trying to decompose.

Local Expansion Decomposition (Section 6). This algorithm, denoted by $\mathcal{A}'$, operates in the *local* setting where there is a graph $G_b = (V, E_b)$ represented by an adjacency list stored in a memory, which was not read by the algorithm. (In our case, G_b will be the connected component C.) Algorithm $\mathcal{A}'$ then takes parameters α_b and ϵ', and a set of edge deletions $D \subseteq E_b$ as inputs. It then gives a decomposition of $G = (V, E_b - D)$ that is similar to that in Theorem 2.3 at a high level; some major differences are:

- It takes $O(\frac{|D|^{1.5+\epsilon'}}{\alpha_b^{3+\epsilon'}} n^{o(1)})$ time, and in particular does not need to read the whole graph G_b.
- It guarantees to outputs some "desired" decomposition only when $\phi(G_b) \geq \alpha_b$.

- In the "desired" decomposition, the expansion of each high-expansion component C is $\phi(C) = \Omega(\alpha_b^{1/\epsilon'})$.

We note the trade-off parameter ϵ' that appears above: when ϵ' is small, the algorithm is fast but has a bad expansion guarantee in the output, and when ϵ' is big, the algorithm is slow but has a good expansion guarantee. We will have to choose this parameter carefully in the end. We state the result about the local decomposition algorithm in more details. (For full details, see Section 6.)

Theorem 2.3. *For any constant $\epsilon' \in (0,1)$, there is an algorithm $\mathcal{A}'$ that can do the following:*

- *$\mathcal{A}'$ is given pointers to G_b,D and α_b stored in a memory: $G_b = (V, E_b)$ is a 3-bounded degree graphs with n nodes represented by an adjacency list. $D \subset E_b$ is a set of edges in G_b. α_b is an expansion parameters. Let $G = (V, E) = (V, E_b - D)$ be the graph that $\mathcal{A}'$ will compute the decomposition on.*
- *Then, in time $O(\frac{|D|^{1.5+\epsilon'}}{\alpha_b^{3+\epsilon'}} n^{o(1)})$, $\mathcal{A}'$ either*
  - *reports failure meaning that $\phi(G_b) < \alpha_b$,*
  - *outputs two graphs $G^s = (V, E^s)$ and $G^d = (V, E^d)$ (except the largest connected component of G^d) where $\{E^s, E^d\}$ is a partition of E.*

  (Note that in the latter case it is still possible that $\phi(G_b) < \alpha_b$, but $\mathcal{A}'$ cannot detect it.)
- *Moreover, if $\phi(G_b) \geq \alpha_b$, then with high probability we have*
  - *$|E^s| = O(|D|/\alpha_b)$, and*
  - *each connected component C of G^d either is a singleton or has high expansion: $\phi(C) = \Omega(\alpha_b^{1/\epsilon'})$.*

Proof Idea (details in 6). Similar to the proof of 2.2, there are two main steps. The first step is to devise a *local* approximation algorithm for a problem called *most-balanced sparse cut with respect to overlapping cuts (MSO cut)*. In this problem, we are given a graph $G = (V, E)$, a target set $A \subset V$ and a parameter α. Then we need to find a α-sparse cut S (i.e. $\phi_G(S) < \alpha$) where $|S| \leq |V - S|$ such that $|S|$ is larger than all α-sparse cuts which are "near" the target set A. ("Nearness" is defined precisely in 6.2). To compare, in the most-balanced sparse cut problem, $|S|$ needs to be larger than *all* α-sparse cuts. By slightly modifying and analyzing the algorithm by Orecchia and Zhu [30] for a related problem called the *local cut improvement problem,* we obtain an approximation algorithm which is local (i.e. its running time depends essentially only on $|A|$, and not $|V|$).

The second step to obtain the decomposition is to find an approximate MSO cut where the target set A is the endpoints of D (together with some additional nodes) and recurse on both sides. To bound the running time, we maintain the same kind of invariant as in the proof of 2.2 throughout the recursion. But, in order to argue that the invariant holds, the analysis is more involved. The main reason is because we compute MSO cuts which have a weaker guarantee, instead of computing most-balanced sparse cuts as in the global expansion decomposition algorithm. However, an important observation is that, when $\phi(G_b) \geq \alpha_b$, any $(\frac{\alpha_b}{2})$-sparse cut S in $G = G_b - D$ must be "near" to D. Intuitively, this is because the expansion of S get halved after deleting edges in D. This justifies why it is enough to find an approximate $(\frac{\alpha_b}{2})$-sparse MSO cut instead of finding an approximate most-balanced sparse cut. As our approximate MSO cut algorithm is local, we can output the decomposition in time essentially independent from the size of G. □

We now finish off our Las Vegas algorithm using the above local decomposition algorithm $\mathcal{A}'$. Let ϵ be a very small constant, and we let $\tau = n^{1/2+\epsilon}$ and $\alpha = n^{\epsilon}$. Recall that we want to maintain a spanning forest of a connected component C undergoing at most τ edge deletions that has expansion α initially, by dividing into phases of τ' updates. In the beginning of each phase, we invoke $\mathcal{A}'$ with $G_b = C$, $\alpha_b = \alpha$, $\epsilon' = \sqrt{\epsilon}$, and D being the set of edges in C deleted so far. Since we start with a component C with expansion at least α, each high-expansion component C' in the resulting decomposition has expansion at least $\alpha' = \Omega(\alpha^{1/\epsilon'})$. Since $|D| \leq \tau$, this algorithm takes time $\lambda = O(\frac{|\tau|^{1.5+\epsilon'}}{\alpha^{3+\epsilon'}} n^{o(1)})$. By plugging these values in Equation (2), we have that the update time of our Las Vegas algorithm is

$$\tilde{O}(\frac{\tau'}{\alpha'} + \frac{\lambda}{\tau'}) = \tilde{O}\left(\frac{\tau'}{\alpha'} + \frac{|\tau|^{1.5+\epsilon'}}{\alpha^{3+\epsilon'}\tau'} n^{o(1)}\right)$$

It is left to pick right parameters to show that the above is $\hat{o}(n^{1/2})$. To this end, we choose $\tau' = \alpha' \times n^{1/2-\epsilon}$, so that the first term in the update time is $n^{1/2-\epsilon}$. The second term then becomes $n^{1/4+O(\epsilon+\epsilon\epsilon'+\epsilon/\epsilon'+\epsilon')}$, which is $\hat{o}(n^{1/2})$ when ϵ is small enough.

Finally, we note that the above algorithm is Las Vegas even though the local algorithm $\mathcal{A}'$ is Monte Carlo for the same reason as we argued before for the case of global algorithm.

Organization

We first introduce relevant notations and definitions in 3. By modifying and combining the known techniques, we present basic tools and give a reduction that proves 2.1 in 4. In 5, we show the global expansion decomposition algorithm which is the backbone of all our dynamic SF algorithms. Then, we extend the idea to show the novel local expansion decomposition algorithm which is a crucial tool for our Las Vegas algorithm, in 6. Finally, we show how all ideas fit together and obtain the Monte Carlo algorithm in 7 and the Las Vegas algorithm in 8.

A contains formal definitions of oblivious and adaptive adversaries. B contains omitted proofs.

3 Preliminaries

In the following, we say that an event e occurs *with high probability* if $\Pr[e] \geq 1 - 1/n^c$ where n is the size of the problem instance depending on the context (in this paper, n is usually a number of nodes in a graph), and c is a fixed constant which can be made arbitrarily large.

Given a graph $G = (V, E)$ and any cut $S \subset V$, we have the following definitions

- the set of cut edges is $\partial_G(S) = \{(u, v) \in E \mid u \in S \text{ and } v \notin S\}$,
- the set of edges induced by the cut S is $E_G(S) = \{(u, v) \in E \mid u, v \in S\}$,
- the cut size of S is $\delta_G(S) = |\partial_G(S)|$,
- the expansion of S is $\phi(S) = \delta(S)/\min\{|S|, |V \setminus S|\}$,
- the volume of S is $vol_G(S) = \sum_{u \in S} \deg(u) = 2|E_G(S)| + \delta_G(S)$, i.e. the number of endpoints of edges in E incident to S.

$V(G) = V$ is the set of nodes of G and $E(G) = E$ is the set of edges of G. The expansion of the graph G is $\phi(G) = \min_{\emptyset \neq S \subset V} \phi(S)$.

For any graph $G = (V, E)$, the *incidence matrix* $B \in \{-1, 0, 1\}^{\binom{|V|}{2} \times |V|}$ of G is defined as following:

$$B_{e,u} = \begin{cases} 1 & \text{if } e = (u, v) \in E \\ -1 & \text{if } e = (v, u) \in E \\ 0 & \text{otherwise.} \end{cases}$$

Let B_u be the u-th column vector of B. We will use the following simple fact.

Proposition 3.1. *For any cut $S \subset V$ and $(i, j) \in \binom{|V|}{2}$, the (i, j)-th entry of the vector $\sum_{u \in S} B_u$ is non-zero iff $(i, j) \in \partial_G(S)$.*

For any graph $G = (V, E)$, let $F \subseteq E$ be a forest in G. We say that an edge $e \in E$ is a *tree edge* if $e \in F$, otherwise e is a *non-tree edge*. For any cut $S \subset V$, let $nt_{G,F}(S) = (E_G(S) \cup \partial_G(S)) \setminus F$ be a set of non-tree edges in G incident to S and let $nt_vol_{G,F}(S) = 2|E_G(S) \setminus F| + |\partial_G(S) \setminus F|$ be the number of endpoints of edges in $E \setminus F$ incident to S.

We say that T is a tree in forest F, or just $T \in F$, if T is a connected component in F. We usually assume that there is a pointer to each tree $T \in F$. For readability, we usually write T instead of a pointer to T in functions of data structures, such as $\mathsf{tree_size}(T)$ in Definition 3.2 and k-$\mathsf{cutset}(T)$ in Definition 4.12.

3.1 Augmented ET Tree

We define the following data structures for convenience. Just by combining the functionality of both ET tree by Henzinger and King [16] with link-cut tree by Sleator and Tarjan [36], that can handle path operation, we have the following:

Definition 3.2 (Augmented ET Tree)**.** *Augmented Euler Tour (ET) tree is a data structure that preprocesses a graph $G = (V, E)$ and a forest $F \subseteq E$ as inputs and handles the following operations:*

- *$\mathsf{insert_tree}(e)$: add edge e into F, given that $e \in E \setminus F$ is a non-tree edge and adding does not cause a cycle in F.*
- *$\mathsf{delete_tree}(e)$: remove edge e from F, given that $e \in F$.*
- *$\mathsf{insert_nontree}(e)$: insert edge e into E, given that $e \notin E$.*
- *$\mathsf{delete_nontree}(e)$: delete edge e from E, given that $e \in E \setminus F$.*
- *$\mathsf{tree_size}(T)$: return $|V(T)|$, given a pointer to T which is a tree in F.*
- *$\mathsf{count_tree}()$: return a number of connected components in F.*
- *$\mathsf{find_tree}(v)$: return a pointer to tree T_v where $v \in T_v$ and T_v is a tree in F.*
- *$\mathsf{sample}(T)$: sample (using fresh random coins) an edge from $nt_{G,F}(S)$ where $S = V(T)$ with probability proportional to its contribution to $nt_vol_{G,F}(S)$, given a pointer to T which is a tree in F. More precisely, an edge from $nt_{G,F}(S) \cap \partial_G(S)$ and $nt_{G,F}(S) \cap E_G(S)$ is sampled with probability $\frac{1}{nt_vol_{G,F}(S)}$ and $\frac{2}{nt_vol_{G,F}(S)}$ respectively.*
- *$\mathsf{list}(T)$: list all edges incident to T, given a pointer to T which is a tree in F.*
- *$\mathsf{path}(u, v, i)$: return the vertex w where the distance from u to w in F is i and w is in a unique path from u to v, given that u and v are in the same tree in F.*

Theorem 3.3. *There is an algorithm for augmented ET tree $\mathcal{E}$ for a graph G with n nodes and m initial edges and a forest F with preprocessing time $O(m)$. $\mathcal{E}$ can handle all operations in $O(\log n)$ time except the following:* tree_size *in time $O(1)$,* count_tree *in time $O(1)$, and* list(T) *in time $O(nt_vol_{G,F}(T))$.*

For completeness, we include the proof in Appendix B.1.

3.2 Definition of Dynamic SF

Definition 3.4 (Dynamic SF)**.** *A dynamic* SF *algorithm $\mathcal{A}$ is given an initial graph to be preprocessed, and then $\mathcal{A}$ must return an initial spanning forest. Then there is an online sequence of edge updates, both insertions and deletions. After each update, $\mathcal{A}$ must returns the list of edges to be added or removed from the previous spanning tree to obtain the new one. We say $\mathcal{A}$ is an incremental/decremental* SF *algorithm if the updates only contain insertions/deletions respectively.*

The time an algorithm uses for preprocessing the initial graph and for updating a new spanning forest is called *preprocessing time* and *update time* respectively.

In this paper, we consider the problem where the update sequence is generated by an adversary.[10] There are two kinds of adversaries: oblivious and adaptive. An *oblivious adversary* initially fixes the whole update sequence, and then shows each update to the algorithm one by one. On the other hand, an *adaptive adversary* generates each update after he sees the previous answer from the algorithm. So each update can depend on all previous answers from the algorithm and all previous updates from the adversaries itself. Adaptive adversaries are obviously as strong as oblivious ones. In Appendix A, we formalize these definitions precisely. By Proposition A.7, we can assume that all adversaries are deterministic.

Throughout the paper, we use the following notation. Given a dynamic SF algorithm $\mathcal{A}$ and an adversary f, let G_0 be the initial graph given by f. Let F_0 be the the initial forest returned from $\mathcal{A}$. For $i \geq 1$, let G_i and F_i be the graph and the forest after the i-th update. More precisely, G_i is obtained by updating G_0 with all updates from f from time 1 to i, and F_i is obtained by updating F_0 with all answers from $\mathcal{A}$ (which describe which tree-edges to be added and removed) from time 1 to i. Since f can be assumed to be deterministic, we have that if $\mathcal{A}$ is deterministic, then G_i and F_i are determined for all $i \geq 0$. On the other hand, if $\mathcal{A}$ is randomized and use a string R as random choices, then G_i and F_i are random variables depending on R for all $i \geq 0$.

Definition 3.5 (Correctness)**.** *For any $p \in [0,1]$ and $T \in \mathbb{N} \cup \{\infty\}$, a dynamic* SF *algorithm $\mathcal{A}$ works correctly against adaptive adversaries with probability p for the first T updates iff, for any fixed adaptive adversary, we have $\Pr_R[F_i$ is a spanning forest of $G_i] \geq p$ for* each *$0 \leq i \leq T$ where R is a random strings used by $\mathcal{A}$ as its random choices. If $T = \infty$, then we omit the phrase "for the first T updates". If $p = 0$, then we say $\mathcal{A}$ works correctly against adaptive adversaries with certainty*

Definition 3.6 (Running Time)**.** *For any $p \in [0,1]$ and $T \in \mathbb{N} \cup \{\infty\}$, a dynamic* SF *algorithm $\mathcal{A}$, that works against adaptive adversaries for the first T updates, has preprocessing time t_p and worst-case update time t_u with probability p iff, for any fixed adaptive adversary, we have $\Pr_R[\mathcal{A}$ preprocesses G_0 and returns F_0 in time at most $t_p] \geq p$ and $\Pr_R[\mathcal{A}$ updates F_{i-1} to be F_i in time at most $t_u] \geq p$ for* each *$1 \leq i \leq T$ where R is a random strings used by $\mathcal{A}$ as its random choices.*

Suppose that F_i is a spanning forest of G_i. If at time $i+1$, a tree-edge $e \in F_i$ is deleted and $F_i \setminus e$ no longer spans $G_{i+1} = G_i \setminus e$. If e' is such that $F_i \cup e' \setminus e$ spans G_{i+1}, then we call that e' is a *replacement edge*. Sometimes we refer to the sequences $G_0, G_1, \ldots$ and $F_0, F_1, \ldots$ as G and F,

[10] This is in contrast to the case where the update sequence is generated by a random process e.g. in [29].

respectively. We call G and F a *dynamic graph* and a *dynamic spanning forest* respectively. We say that $\mathcal{A}$ *maintains* F *in* G. G has m *initial edges* if $E(G_0) \leq m$, and G has at most m edges if $E(G_i) \leq m$ for all i.

4 Basic Tools

In this section, we present some useful tools for our dynamic SF algorithms. All the tools presented in this section are obtained by combining or modifying the existing techniques. This include 2-dimensional ET Tree in Section 4.1, cut recovery tree in Section 4.2, and a reduction which greatly simplifies our goal in Section 4.3.

4.1 2-dimensional ET Tree

The following data structure is useful for both of our Monte Carlo and Las Vegas dynamic SF algorithms. It will be used to help maintaining a spanning forest in a graph with few non-tree edges.

Definition 4.1 (2-dim ET Tree)**.** *2-dimensional (2-dim) ET tree is a data structure that preprocesses a graph $G = (V, E)$ and a forest $F \subseteq E$ as inputs and handles the following operations:*

- *insert_tree*(e)*,* *delete_tree*(e)*,* *insert_nontree*(e)*,* *delete_nontree*(e) *and* *find_tree*(u)*: same definitions as in Definition 3.2.*
- *find_edge*(T)*: return an edge $(u, v) \in \partial_G(V(T))$ if exists, given a pointer to T which is a tree in F.*
- *find_edge*(T_1, T_2)[11]*: return an edge (u, v) where $u \in V(T_1)$ and $v \in V(T_2)$ if exists, given pointers to T_1 and T_2 which are both trees in F.*

We say that the number of non-tree edges of the underlying graph $G = (V, E)$ is bounded by k, if throughout all the update operations, we have $|E - F| \leq k$.

Theorem 4.2. *There is a deterministic data structure for 2-dim ET tree denoted by $\mathcal{E}^{2d}$ that preprocesses G, F and a number k as inputs where G is a graph with n nodes, F is a forest in G, and the number of non-tree edges of G is always bounded by k. $\mathcal{E}^{2d}$ preprocesses the input in time $O(n + k)$, and can handle all operations in $O(\sqrt{k} \log k)$ time.*

There are several data structures based on topological tree [12] that can be used to prove a slightly weaker version of Theorem 4.2. These data structures include 2-dimensional topological tree by Frederickson [12], and a generalization of top tree by Alstrup et al. [2, Section 5] or by Thorup [40, Section 5.3]. They handle all operations of 2-dim ET tree in time $\tilde{O}(\sqrt{m})$ where m is the number of *all edges*. However, in our application, we need to handle the operations in time $\tilde{O}(\sqrt{k})$ where k is the number of *non-tree edges* but k might be much less than m.

We believe that, with some careful modification, these topological-tree-based data structures can be used to prove Theorem 4.2. Hence, this theorem may be considered as a folklore. But, for completeness, we instead show an implementation that is based on ET tree instead which is simpler than topological tree and then prove Theorem 4.2. Again, even this ET-tree-based implementation might be considered as a folklore.

[11] In the subsequent sections, the operation find_edge(T_1, T_2) will not be used. But we state it because it can be useful for other applications.

For the rest of this section, we will show how to construct 2-dim ET tree data structure that cannot handle find_tree operation. But it is easy how to augment it so that it can handle find_tree, for example by combining with augmented ET tree.

4.1.1 Reduction to a Constant Number of Connected Components

We first show the reduction that allows us to assume that there are only $O(1)$ number of connected components in F, i.e. there are only $O(1)$ trees in F. Recall that we call each connected component T in F as a tree $T \in F$.

Lemma 4.3. *Suppose there is a 2-dim ET tree data structure $\mathcal{D}$ that preprocesses any (G, F) in time $t_p(n, k)$ where G is a graph with n nodes and F is a forest in G, and the number of non-tree edges of G is bounded by k. Moreover, all operations of $\mathcal{D}$ can done in $t_u(n, k)$ time as long as the number of connected component in F is bounded by a constant. Then there is a 2-dim ET tree data structure $\mathcal{E}$ that preprocesses any (G, F) in time $O(n + k + t_p(n, k))$ where G is a graph with n nodes and F is a forest in G, and the number of non-tree edges of G is bounded by k. Moreover, all operations of $\mathcal{E}$ can done in $O(t_u(n, k))$ time.*

Proof. We will show how get obtain $\mathcal{E}$ given that we have $\mathcal{D}$. Given $G = (V, E)$ and $F \subset E$, $\mathcal{E}$ preprocesses (G, F) as follows. First, we construct $G' = (V', E')$ and F'. Let r be an additional node and let $V' = V \cup \{r\}$. The invariant of F' is such that for each connected component T in F, there is a tree edge $(r, u_T) \in F'$ where $u_T \in V(T)$ is some node in T. Let $E' = E \cup F'$. Then, we preprocess (G', F') using $\mathcal{D}$. The total time is $O(n + k) + t_p(n, k)$. Note that F' has one connected component.

Given the operation insert_tree(e), delete_tree(e), insert_nontree(e), delete_nontree(e) for updating G and F, it is clear how to update G' and F' such that the invariant is maintained: for each $T \in F$, there is an edge $(r, u_T) \in F'$ where $u_T \in V(T)$ using a constant number of update operations to G' and F' while making sure that F' has at most $O(1)$ many connected components.

Given a query find_edge(T) where $T \in F$, we give the following update to G' and F': delete_tree(r, u_T), find_edge(T), insert_tree(r, u_T). Similarly, given a query find_edge(T_1, T_2) where $T_1, T_2 \in F$, we update G' and F' as follows: delete_tree(r, u_{T_1}), delete_tree(r, u_{T_2}), find_edge(T_1, T_2), insert_tree(r, u_{T_1}), insert_tree(r, u_{T_2}). It is clear also answers from $\mathcal{D}$ maintain (G', F') are correct and takes at most $O(1) \times t_u(n, k) = O(t_u(n, k))$ time. This concludes the reduction. □

4.1.2 List Intersection Oracle

In this section, we show the main data structure which is used inside 2-dim ET tree. For any list of numbers $L = (e_1, \ldots, e_s)$, we denote $|L| = s$ the length of the list. We say that a number e *occurs* in L for c times if $|\{j \mid e = e_j \in L\}| = c$. Let $\mathcal{L} = \{L_i\}_i$ be a collection of lists. A number e occurs in $\mathcal{L}$ for c times if $\sum_{L \in \mathcal{L}} |\{j \mid e = e_j \in L\}| = c$.

Definition 4.4 (List intersection oracle)**.** List intersection oracle *is a data structure that preprocesses a collection of lists of numbers $\mathcal{L} = \{L_i\}_i$ as inputs and handles the following operations:*

- *create*(e) : *return a pointer to a new list $L = \{e\}$, given that e is a number.*
- *merge*(L_1, L_2)*: concatenate two lists into a new list $L = L_1 L_2$, given pointers to $L_1, L_2 \in \mathcal{L}$.*
- *split*(L, i)*: split $L = (e_1, \ldots, e_s)$ into two lists L_1 and L_2 where $L_1 = (e_1, \ldots, e_i)$ and $L_2 = (e_{i+1}, \ldots, e_s)$, given a pointer to $L \in \mathcal{L}$ and $1 \leq i \leq s$.*

- *intersection*(L_1, L_2): *return an element* e *where* $e \in L_1 \cap L_2$ *if* v *exists (otherwise return* $\perp$*), given two pointers two* $L_1, L_2 \in \mathcal{L}$.

Theorem 4.5. *For any parameter* k*, there is a list intersection oracle* $\mathcal{D}$ *that preprocess a collection of lists* $\mathcal{L} = \{L_1\}_i$ *as inputs in time* $O(k)$*. Suppose that 1) the total length of lists in* $\mathcal{L}$ *is bounded by* $\sum_i |L_i| \leq k$*, 2) the number of lists in* $\mathcal{L}$ *is bounded by some constant, and 3) each number occurs in* $\mathcal{L}$ *at most a constant number of times. Then* $\mathcal{D}$ *can handle* merge *and* split *in* $O(\sqrt{k} \log k)$ *time, and can handle* create *and* intersection *in* $O(1)$ *time.*

We note again that all the ideas of this algorithm have been presented already by Alstrup et al. [2, Section 5] and by Thorup [40, Section 5.3]. Since they did not explicitly define a list intersection oracle as a problem, we include it for completeness in Appendix B.2.1.

4.1.3 Non-tree Euler Tour

The goal of this section is the following:

Theorem 4.6. *There is a 2-dim ET tree data structure* $\mathcal{D}$ *that preprocesses any* (G, F) *in time* $O(n+k)$ *where* G *is a graph with* n *nodes and* F *is a forest in* G*, and the number of non-tree edges of* G *is bounded by* k*. Moreover, all operations of* $\mathcal{D}$ *can done in* $O(\sqrt{k} \log k)$ *time as long as the number of connected component in* F *is bounded by a constant.*

By Lemma 4.3, this immediately implies Theorem 4.2.

The key to this result is the definition of *non-tree Euler Tour* of a tree, which is a slight modification of Euler Tour defined in [16]. While Euler tour is an order list of nodes in a tree, non-tree Euler Tour is a list of non-tree edges incident to nodes in a tree. The construction is almost the same.

Throughout this section, there are only $O(1)$ number of trees $T \in F$. For each tree $T \in F$, we root T at arbitrary node r_T. For each node $u \in T$, let $p_T(u)$ be the parent of u in T, and we fix an arbitrary ordering to edges incident to u.

Definition 4.7 (Non-tree Euler Tour)**.** *For any graph* $G = (V, E)$*, forest* $F \subset E$*, and tree* $T \in F$*, a* non-tree Euler tour of T *in* G*, denoted by* $\mathsf{NET}_G(T)$*, is an ordered list of non-tree edges in* G *that is incident to vertices in* T*. The precise ordering in* $\mathsf{NET}_G(T)$ *is defined by the output of* $\mathsf{ListNET}(G, F, r_T)$ *where* r_T *is the root of* T*. See Algorithm 4.1.* Non-tree Euler tour of F *in* G*, denoted by* $\mathsf{NET}_G(F) = \{\mathsf{NET}_G(T)\}_{T \in F}$*, is a collection of the tours of trees* $T \in F$*.*

1. For each edge (u, v) incident to u:
   (a) If $v = p_T(u)$, CONTINUE.
   (b) If $(u, v) \notin F$, then OUTPUT $e = \{u, v\}$.
   (c) If $(u, v) \in F$, then $\mathsf{ListNET}(G, F, v)$.

Algorithm 4.1: $\mathsf{ListNET}(G, F, v)$ generates a list of non-tree edges of G incident to vertices in a subtree in F rooted at v.

Next, we observe some easy properties of non-tree Euler tour which are similarly shown for Euler tour. The number of occurrence of an edge e in $\mathsf{NET}_G(F)$ is $o(e) = \sum_{T \in F} |\{i \mid e = e_i \in \mathsf{NET}_G(T)\}|$.

Proposition 4.8. *Each non-tree edge of G occurs in $\mathsf{NET}_G(F)$ exactly two times.*

Proposition 4.9. *For any $T_1, T_2 \in F$, there is a non-tree edge (u, v) where $u \in V(T_1)$ and $v \in V(T_2)$ iff $\mathsf{NET}_G(T_1) \cap \mathsf{NET}_G(T_2) \neq \emptyset$.*

Proposition 4.10. *Let $G = (V, E)$ be a graph and $F \subset E$ be a forest in G, and $|E - F| = k$. Then $\sum_{T \in F} |\mathsf{NET}(T)| = 2k$.*

Now, we describe how to use $\mathsf{NET}_G(F)$ for 2-dimenstion ET tree. Let $\mathcal{O}$ be list intersection oracle from Theorem 4.5 with parameter $2k$. By Proposition 4.10, Proposition 4.8, and the fact that $O(1)$ number of trees $T \in F$, the conditions in Theorem 4.5 are satisfied. And hence $\mathcal{O}$ can handle merge and split in $O(\sqrt{k} \log k)$ time, and can handle create and intersection in $O(1)$ time. We have the following:

Proposition 4.11. *Let $G = (V, E)$ be a graph and $F \subset E$ be a forest in G. Suppose that there is a list intersection oracle $\mathcal{O}$ that has preprocessed $\mathsf{NET}_G(F) = \{\mathsf{NET}_G(T)\}_{T \in F}$.*

- *For any $T_1, T_2 \in F$, the answer of find_edge(T_1, T_2) and be obtained from intersection$(\mathsf{NET}_G(T_1), \mathsf{NET}_G(T_2))$.*
- *Let G' and F' be obtained from G and F by either inserting or deleting one tree edge or non-tree edge. Then $\mathsf{NET}_{G'}(F')$ can be obtained from by a constant number of create, merge and split on $\mathcal{D}$.*

Now, we can prove Theorem 4.6:

Proof of Theorem 4.6. Given G and F, we can constructing the non-tree Euler tour $\mathsf{NET}_G(F)$ in time $O(n + k)$. Then we give $\mathsf{NET}_G(F)$ to the list intersection oracle $\mathcal{O}$ as inputs, and $\mathcal{O}$ takes $O(k)$ times to preprocess. The worst-case time of any operation handled by $\mathcal{O}$ is $O(\sqrt{k} \log k)$ by Theorem 4.5. By Proposition 4.11, we can do insert_tree(e), delete_tree(e), insert_nontree(e), delete_nontree(e), and find_edge(T_1, T_2) in time $O(\sqrt{k} \log k)$. To answer find_edge(T), we can just query find_edge(T, T') for all $T' \neq T$. □

4.2 Cut Recovery Tree

The following data structure is useful for our Monte Carlo dynamic SF algorithm.

Definition 4.12 (Cut Recovery Tree)**.** *Cut recovery tree is a data structure that preprocesses a graph $G = (V, E)$, a forest $F \subseteq E$ and a parameter k as inputs and handles the following operations:*

- *insert_tree(e), delete_tree(e), insert_nontree(e) and delete_nontree(e): same definitions as in Definition 3.2.*
- *k-cutset(T): return all edges in $\partial_G(V(T))$, given a pointer to T which is a tree in F and given that $\delta_G(V(T)) \leq k$.*

Theorem 4.13. *There is a deterministic data structure for cut recovery tree denoted by $\mathcal{D}$ that are given G, F and a parameter k as inputs where $G = (V, E)$ is a graph with n nodes and m initial edges, and $F \subseteq E$ is a forest in G. $\mathcal{D}$ can preprocess the inputs in time $O(m\beta \log n)$ time where $\beta = 2^{O(\log \log n)^3} = n^{o(1)}$. And then $\mathcal{D}$ takes*

- *$O(k\beta \log^2 n)$ time to handle insert_tree and delete_tree,*
- *$O(\beta \log^2 n)$ time to handle insert_nontree and delete_nontree, and*

- $O(k\beta \log^2 n)$ *time to answer* k-*cutset*.

This theorem is proved by maintaining a version of ET tree where each node of the ET tree contains a k-sparse recovery sketch. This approach is already used by Kapron, King and Mountjoy [23] to obtain their data structure called *cut set data structure*. But in [23], each node of the ET tree contains instead contains a ℓ_0-sampling sketch (which is essentially a 1-sparse recovery sketch combining with sampling). Here, we generalize the same idea to k-sparse recovery sketch.

A nice aspect about both cut set data structure and cut recovery tree is that they are obtained by combining ideas from two different communities: 1) linear sketches from streaming literature, see Section 4.2.1, and 2) a version of ET tree which is a classic dynamic data structure, see Section 4.2.2. We show how to combine the two and prove Theorem 4.13 in Section 4.2.3.

4.2.1 Sparse Recovery

In this section, we show the following about sparse recovery. We say that a vector v is k-*sparse* if there are at most k non-zero entries in v.

Theorem 4.14. *There are algorithm $\mathcal{A}_{recover}$, and algorithm $\mathcal{A}_{construct}$ such that, given any number n and k, they can do the following. Let $\beta_1 = 2^{O(\log\log n)^3} = n^{o(1)}$, $\beta_2 = O(polylogn)$, $d = O(k\beta_1 \log n)$ and $s = O(\beta_1 \log n)$.*

- *There is a matrix $\Phi \in \{0,1\}^{d\times n}$ where each column of Φ is s-sparse that can be computed by $\mathcal{A}_{construct}$ in the following sense: given an index $i \in [n]$, $\mathcal{A}_{construct}$ can return a list of non-zero entries in the i-th of column of Φ, i.e. $(j, \Phi_{j,i})$ where $\Phi_{j,i} \neq 0$ for each $j \in [d]$, in time $s\beta_2$ time.*
- *For any k-sparse vectors $v \in [-M, M]^n$ where $M = poly(n)$, given Φv, $\mathcal{A}_{recover}$ can return a list of non-zero entries in v, i.e. the list of (i, v_i) where $v_i \neq 0$ for each $i \in [n]$, in time $O(d \log n)$.*

This theorem is essentially proved by Berinde et al. [5, Section 5]. However, in [5], they did not explicitly state how the matrix Φ can be constructed and the running time of the construction. Here, we make it explicit by defining the algorithm $\mathcal{A}_{construct}$. To do this, we observe that the construction of [5] uses an explicit *unbalanced expander* from Capalbo et al. [9]. The explicitness of their expanders directly allows us to obtain $\mathcal{A}_{construct}$. We also give the proof of this theorem for completeness in Appendix B.2.2.

4.2.2 d-word ET Tree

Next, we need the following data structure which is a variant of ET tree from Henzinger and King's paper [16] so that it can maintain a sum of a large vector associated with each node. This data structure is also used by Kapron King and Mountjoy [23] to obtain a data structure call cut set data structure. Since the definition is implicit there, we give a precise definition here:

Definition 4.15 (d-word ET tree)**.** *For any parameter d, a d-word ET tree is a data structure that preprocesses a forest F with a set V of nodes and a vector $x_u \in [-M, M]^d$ where $M = poly(|V|)$ for each $u \in V$ as inputs and handles the following operations:*

- *insert_tree*(e)*: add edge e into F, given that $e \notin F$ and adding does not cause a cycle in F.*
- *delete_tree*(e)*: remove edge e from F, given that $e \in F$.*

- *sum_vec*(T): *return* $\sum_{u\in T} x_u$, *given that* T *is a tree in* F.
- *update_vec*(u, i, α): *set* $x_u[i] = \alpha$, *given that* $u \in V$, $i \in [d]$ *and* $\alpha \in [-M, M]$.

Theorem 4.16. *For any parameter* d, *there is an algorithm for* d*-word ET tree that preprocesses a forest* F *with a set* V *of* n *nodes and a set of vectors* $x = \{x_u\}_{u\in V}$ *which takes*

- $O(nnz(x) \log n)$ *preprocessing time where* $nnz(x) = \sum_{u\in V} |x_u|_0$ *is the total number of non-zero entries of each vector in* $\{x_u\}_{u\in V}$,
- $O(d \log n)$ *time to handle* *insert_tree* *and* *delete_tree*,
- $O(d)$ *time to return* *sum_vec*, *and*
- $O(\log n)$ *time to handle* *update_vec*(u, i, α).

It is important that preprocessing time is $\tilde{O}(nnz(x))$ to obtain a cut recovery tree with fast preprocessing time. Note that the total number of entries of x_u over all $u \in V$ is $n \times d$. But $nnz(x)$ can be much smaller. For completeness, we give a proof in Appendix B.2.3.

4.2.3 Put Them Together

Theorem 4.13 is obtained by combining the above two results:

Proof of Theorem 4.13. Let Φ^* be the matrix of size $d \times \binom{n}{2}$ from Theorem 4.14 with parameter $\binom{n}{2}$ and k such that each column of Φ^* is s-sparse, where $d = O(k\beta_1 \log n)$, $s = O(\beta_1 \log n)$ and $\beta_1 = 2^{O(\log\log n)^3}$. We will not actually construct the whole Φ^* at the beginning. Instead, let Φ be a zero matrix of size $d \times \binom{n}{2}$. We will gradually compute some entries of $\Phi^*$and copy it to the corresponding entry in Φ. For each $(u, v) \in \binom{n}{2}$, let $\Phi_{(u,v)}$,$\Phi^*_{(u,v)} \in \{0, 1\}^d$ denote the (u, v)-th column of Φ and Φ^* respectively.

Let E_0 be the set of initial edges of G. At any time, let E_I be the set of edges have been inserted into G and let $B \in \{0, 1\}^{\binom{n}{2}\times n}$ be the incidence matrix of G. We have the following invariants: 1) for each $(u, v) \in E_0 \cup E_I$, $\Phi_{(u,v)} = {\Phi^*}_{(u,v)}$, otherwise $\Phi_{(u,v)} = 0$, and 2) for each $u \in V$, we have a vector $x_u = \Phi B_u$ of dimension d, where $B_u \in \{0, 1\}^{\binom{n}{2}}$ be the u-th column of B. It follows from the invariants that, for any $S \subset V$, we have $\Phi(\sum_{u\in S} B_u) = \Phi^*(\sum_{u\in S} B_u)$ because if the (u, v)-th entry of $\sum_{u\in S} B_u$ is non-zero, then $(u, v) \in E_0 \cup E_I$.

To preprocess, we do the following. We assume that it takes constant time to allocate space for a zero matrix Φ. Then, to validate the invariants, we first set $\Phi_{(u,v)} = {\Phi^*}_{(u,v)}$ for each initial edge $(u, v) \in E_0$. This takes in total $ms\beta_2$ time because all non-zero entries of each column of Φ^* can be identified in $s\beta_2$ time by Theorem 4.14 where $\beta_2 = O(\text{polylog}n)$. Second, we compute $x_u = \Phi B_u$. Since each column of Φ^* is s-sparse, so is Φ. As, B_u is $\deg(u)$-sparse, this also takes in total $\sum_{u\in V} O(\deg(u)s) = O(ms)$ time. Finally, let $\mathcal{E}$ be a d-word ET tree from Theorem 4.16. We use $\mathcal{E}$ to preprocess F and $x = \{x_u\}_{u\in V}$ taking time $O(nnz(x) \log n) = O(\sum_{u\in V} \deg(u) \log n) = O(m \log n)$. The total preprocessing time is $O(ms\beta_2 + ms + m \log n) = O(m\beta_1\beta_2 \log n)$.

To handle insert_tree and delete_tree, we just update $\mathcal{E}$ accordingly. So, by Theorem 4.16, the tree-edge update time is $O(d \log n) = O(k\beta_1 \log^2 n)$. To handle insert_nontree and delete_nontree, if $e = (u, v)$ is inserted and $e \notin E_0 \cup E_I$, then, to maintain the first invariant, we set $\Phi_{(u,v)} = {\Phi^*}_{(u,v)}$ in time $s\beta_2$. Next, to maintain the second invariant, we want to maintain that $x_u = \Phi B_u$ for all $u \in V$. When we update non-tree edge $e = (u, v)$, one entry of both B_u and B_v are updated. Since each column of Φ is s-sparse, s entries of x_u and x_v are updated. Each call for update_vec takes

$O(\log n)$ time, so this takes $O(s \log n)$. In total, the non-tree-edge update time is $O(s\beta_2 + s\log n) = O(\beta_1\beta_2 \log n)$.

To answer k-$\mathsf{cutset}(T)$, we have that for any fixed T, we can construct $\partial_G(V(T))$ from $\sum_{u\in T} x_u$ in time $O(d \log n)$. Because 1) by calling $\mathsf{sum_vec}(T)$, we get $\sum_{u\in T} x_u$ in time $O(d)$, 2) $\sum_{u\in T} x_u = \sum_{u\in T} \Phi B_u = \Phi(\sum_{u\in T} B_u) = \Phi^*(\sum_{u\in S} B_u)$, 3) $\partial_G(V(T))$ corresponds to the non-zero entries of $\sum_{u\in T} B_u$ by Proposition 3.1, and 4) by Theorem 4.14, for any k-sparse $v \in [-poly(n), poly(n)]^{\binom{n}{2}}$, we can construct v from from $\Phi^* v$ in time $O(d \log n)$. Since $\delta_G(V(T)) \leq k$, we can construct $\partial_G(V(T))$ in time $O(d \log n) = O(k\beta_1 \log^2 n)$.

By setting the parameter $\beta = \beta_1 \times \beta_2$, this completes the proof. □

Remark 4.1. We note that β is a factor depending on product of the degree and explicitness of the best known explicit construction of unbalanced expanders. Therefore, a better construction of expanders would immediately speed up our data structure.

4.3 Classic Reductions

In this section, we show that, roughly, to get a dynamic SF with $O(n^{1/2-\epsilon+o(1)})$ worst-case update time for an infinite-length update sequence, it is enough to show a dynamic SF with the same update time and $O(n^{1+o(1)})$ preprocessing time where the underlying graph has maximum degree 3 and the length of the update sequence is just $n^{1/2+\epsilon}$.

This reduction also helps the algorithms in [23, 13] that can handle only polynomial-length update sequences to be able handle infinite-length sequences.

Recall Definition 3.5 and Definition 3.6 for formal definitions of "working against adversaries". We also need the following definition:

Definition 4.17. *A dynamic algorithm runs* on Δ-bounded degree graphs *if throughout the whole sequence of updates, all fnodes in the underlying graph always has maximum degree at most* Δ.

Theorem 4.18. *For any n, p_c, p_t and T, suppose there is a dynamic SF algorithm $\mathcal{A}$ for 3-bounded-degree graphs with n nodes and $\Theta(n)$ edges that, for the first T updates, works correctly against adaptive adversaries with probability $1-p_c$ and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and worst-case update time $t_u(n, p_c, p_t, T)$ with probability $1-p_t$. Suppose that for any n_1, n_2, $t_p(n_1, p_c, p_t, T) + t_p(n_2, p_c, p_t, T) \leq t_p(n_1+n_2, p_c, p_t, T)$, then there is a dynamic SF algorithm $\mathcal{B}$ for any graph with n nodes that works correctly against adaptive adversaries with probability $1 - p_c \cdot poly(n)$ and $\mathcal{B}$ has preprocessing time $O(t_p(O(m), p_c, p_t, T) \log n)$ where m is the number of initial edges and worst-case update time $O(t_u(O(n), p_c, p_t, T) \log n + t_p(O(n), p_c, p_t, T)/T + \log^2 n)$ with probability $1 - p_t \cdot poly(n)$.*

Note that if $\mathcal{A}$ works correctly with certainty ($p_c = 0$) or with high probability, then so does $\mathcal{B}$. Similarly,if $\mathcal{A}$ guarantees the worst-case update time with certainty ($p_t = 0$) or with high probability, then so does $\mathcal{B}$.

For convenience, we state the reduction that follows immediately from Theorem 4.18 and will be used in Section 7 to obtain our Monte Carlo algorithm.

Corollary 4.19. *For any constant ϵ, suppose there is a dynamic SF algorithm $\mathcal{A}$ for 3-bounded-degree graphs with n nodes and $\Theta(n)$ edges that, for the first T updates, works correctly against adaptive adversaries with probability $1-p_c$ and $\mathcal{A}$ has preprocessing time $t_p(n, p_c) = O(n^{1+o(1)} \log \frac{1}{p_c})$ and worst-case update time $t_u(n, p_c, T)$ where $t_u(n, p_c, n^{1/2+\epsilon}) = O(n^{1/2-\epsilon+o(1)} \log \frac{1}{p_c})$. Then there is a dynamic SF algorithm $\mathcal{B}$ for any graph with n nodes that works correctly against adaptive adversaries with high probability and $\mathcal{B}$ has preprocessing time $O(m^{1+o(1)})$ where m is the number of initial edges and worst-case update time $O(n^{1/2-\epsilon+o(1)})$.*

Proof. We note that the function $t_p(n, p_c)$ does not depend on T. We set $T = n^{1/2+\epsilon}$. By Theorem 4.18, there is dynamic SF algorithm $\mathcal{B}$ for any graph with n nodes that works correctly against adaptive adversaries with probability $1 - p_c \cdot \mathrm{poly}(n)$ and $\mathcal{B}$ has preprocessing time $O(t_p(O(m), p_c) \log n) = O(m^{1+o(1)} \log \frac{1}{p_c})$ where m is the number of initial edges and worst-case update time $O(t_u(O(n), p_c, T) \log n + t_p(O(n), p_c)/T + \log^2 n) = O(n^{1/2-\epsilon+o(1)} \log \frac{1}{p_c} + \frac{n^{1+o(1)} \log \frac{1}{p_c}}{n^{1/2+\epsilon}} + \log^2 n) = O(n^{1/2-\epsilon+o(1)} \log \frac{1}{p_c})$. By choosing $p_c = 1/n^k$ where k is an arbitrarily large constant, we are done. □

Next, we state the similar reduction that will be used in Section 8 to get our Las Vegas algorithm.

Corollary 4.20. *For any n and constant ϵ, suppose there is a dynamic SF algorithm $\mathcal{A}$ for 3-bounded-degree graphs with n nodes and $\Theta(n)$ edges that, for the first T updates, works correctly against adaptive adversaries with certainty and $\mathcal{A}$ has, with probability $1 - p_t$, preprocessing time $t_p(n, p_t) = O(n^{1+o(1)} \log \frac{1}{p_t})$ and worst-case update time $t_u(n, p_t, T)$ where $t_u(n, p_t, n^{1/2+\epsilon}) = O(n^{1/2-\epsilon+o(1)} \log \frac{1}{p_t})$. Then there is a dynamic SF algorithm $\mathcal{B}$ for any graph with n nodes that works correctly against adaptive adversaries with certainty and $\mathcal{B}$ has preprocessing time $O(m^{1+o(1)})$ where m is the number of initial edges and worst-case update time $O(n^{1/2-\epsilon+o(1)})$ with high probability.*

Corollary 4.20 has almost the same proof as Corollary 4.19 so we omit it. The rest of this section is devoted for proving Theorem 4.18.

4.3.1 Overview of the Sequence of Reductions

In this section, we will state the sequence of reductions that we will use for proving Theorem 4.18. For each reduction, we give a high level overview. Then we show how to combine them. For convenience, we define dynamic SF algorithms with parameters as follows:

Definition 4.21. *(n, p_c, p_t, T)-algorithms are dynamic SF algorithms for a graph with n nodes that, for the first T updates, works correctly against adaptive adversaries with probability $1 - p_c$ and the preprocessing and update time is guaranteed with probability $1 - p_t$. $(n, p_c, p_t\infty)$-algorithms give the same guarantee for correctness and running time for any length of updates.*

The following property called "stability" of the algorithm is useful for us. Note that this definition is not the same definition as in [11, Definition 3.3.1] as we will discuss again before the statement of Lemma 4.27.

Definition 4.22 (Stability)**.** *A dynamic SF algorithm $\mathcal{A}$ is* stable *if the spanning forest maintained by $\mathcal{A}$ never changes except when a tree-edge e is deleted. If a tree-edge e is deleted from a spanning forest F, then $\mathcal{A}$ either do not update F or add a replacement edge e' into F. That is, let F' be a new spanning forest updated by $\mathcal{A}$. Then, either $F' = F \setminus \{e\}$ or $F' = F \cup \{e'\} \setminus \{e\}$.*[12]

An (n, p_c, p_t, T)-algorithm which is stable is called a stable (n, p_c, p_t, T)-algorithm. The first reduction shows that we can assume that the algorithm is stable. The technique for proving this is implicit in [16, Section 3.1].

[12] All known dynamic spanning tree algorithms in the literature [12, 23, 19, 16, 24] are stable according to this definition.

Lemma 4.23. *Suppose there is an (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph with preprocessing time t_p and update time t_u. Then there is a stable $(n, O(Tp_c), p_t, T)$-algorithm $\mathcal{B}$ for the same graph with preprocessing time $O(t_p + m_0)$ where m_0 is the number of initial edges, and update time $O(t_u \log n + \log^2 n)$.*

Next, the reduction below shows how to get an algorithm that can handle long update sequences from an algorithm that handles only short sequences. The technique is standard, i.e. dividing the long update sequence into phases, and maintaining two concurrent data structures.

Lemma 4.24. *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$ with preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then, for any $U \geq T$, there is a stable $(n, O(Up_c), O(p_t), U)$-algorithm $\mathcal{B}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$ with preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) + t_p(n, p_c, p_t, T)/T + \log^2 n)$.*

The next reduction shows how to assume that the graph is 3-bounded-degree. The reduction is implicit in [12, 24], so the technique is standard. Given a graph G, we maintain a graph G' with max degree 3 by "splitting" each node in G into a path in G' of length equals to its degree. We assign weight 0 to edges in those paths, and 1 to original edges. Then we maintain a 2-weight minimum spanning forest in G'. To do this, we use the reduction from k-weight minimum spanning forest to spanning forest from [16].

Lemma 4.25. *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$, and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then there is a stable $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{B}$ for any graph with where the number of edges is always between $n/4$ and $2n$ with preprocessing time $O(t_p(O(n), p_c, p_t, T))$ and update time $O(t_u(O(n), p_c, p_t, T) + \log^2 n)$.*

The following shows how to obtain an algorithm whose running time depends only on the number of edges, and not depends on the number n of nodes. It is useful when the number of edges is *much less than* the number of nodes. The technique for this reduction is also standard, i.e. dividing the long update sequence into phases, and maintaining two concurrent data structures.

Lemma 4.26. *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph where the number of edges is always between $n/4$ and $2n$, and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then there is a stable $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{B}$ for any graph, initially represented as a list of its edges[13], with n nodes where the number of edges is always at most $2n$, and $\mathcal{B}$ has preprocessing time $O(t_p(O(m_0), p_c, p_t))$ where m_0 is a number of initial edges, and update time $O(t_u(O(m), p_c, p_t) + \log^2 n)$ where m is a number of edges when update.*

The next reduction shows, given an algorithm with running time depending only on the number of edges, how to get an algorithm with update time depending only on the number of nodes in even dense graphs. This reduction uses the general sparsification technique introduced in [11]. However, it is not immediate that we can apply their framework because their notion of *stability* is too strong for us.

To be more precise, a *stable certificate* according to [11, Definition 3.3.1] must be uniquely determined once the graph is fixed. However, this is not the case for a spanning forest maintained

[13]Note that it is possible that the preprocessing time of $\mathcal{B}$ is $O(t_p(O(m_0), p_c, p_t)) = o(n)$, so the algorithm $\mathcal{B}$ needs that the initial graph G *is represented as a list of its edges* i.e. G is represented as a list $((u_1, v_1), (u_2, v_2) \ldots, (u_m, v_m))$ so that all edges of G can be listed in time $O(m)$ even when $m = o(n)$.

by a randomized algorithm, because its edges may depend on the random choices of the algorithm. So we need to use the weaker notion of stability as defined in Definition 4.22 and prove that this framework still goes through.

Lemma 4.27. *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a graph, initially represented as a list of its edges, with n nodes where the number of edges is always at most $2n$, and $\mathcal{A}$ has preprocessing time $t_p(m_0, p_c, p_t, T)$ where m_0 is a number of initial edges and update time $t_u(m, p_c, p_t, T)$ where m is a number of edges when update. Suppose that for any n_1, n_2, $t_p(m_1, p_c, p_t, T)+t_p(m_2, p_c, p_t, T) \leq t_p(m_1 + m_2, p_c, p_t, T)$, then there is a stable $(n, O(np_c), O(np_t), T)$-algorithm $\mathcal{B}$ for any graph with n nodes with preprocessing time $O(t_p(m_0, p_c, p_t, T) \log n)$ where m_0 is a number of initial edges and update time $O(t_u(O(n), p_c, p_t, T))$.*

Observe that many reductions above, given an (n, p_c, p_t, T)-algorithm $\mathcal{A}$, we obtain the (n, p'_c, p'_t, T')-algorithm $\mathcal{B}$ where $p'_c = \Omega(T'p_c)$. That is the failure probability for correctness is always increased by a factor of T', which is the length of update sequence that the resulting algorithm $\mathcal{B}$ can handle. The following reduction shows that, if $T \geq n^2$, how to obtain an algorithm that can handle update sequence of *any length*, while the failure probability does not increase too much. This kind of reduction has been shown for spanners in [1]. However, the reduction from [1] does not work because the structure of spanning forests is more restricted than spanners.

Lemma 4.28. *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph with n nodes with preprocessing time $t_p(m_0, p, p_t, T)$ where m_0 is a number of initial edges and update time $t_u(n, p_c, p_t, T)$. If $T \geq n^2$, then there is an $(n, O(Tp_c), O(p_t), \infty)$-algorithm $\mathcal{B}$ for any graph with n nodes with preprocessing time $O(t_p(m_0, p_c, p_t, T))$ where m_0 is a number of initial edges and update time $O(t_u(n, p_c, p_t, T))$.*

Remark 4.2. The spanning forest maintained by the algorithm $\mathcal{B}$ in Lemma 4.28 is not stable.

With all reductions above, we conclude the following which is equivalent to Theorem 4.18.

Theorem 4.29 (Restatement of Theorem 4.18)**.** *Suppose there is an (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a 3-bounded-degree graph where the number of edges is always between $n/2$ and $2n$, and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Suppose that for any n_1, n_2, $t_p(n_1, p_c, p_t, T)+t_p(n_2, p_c, p_t, T) \leq t_p(n_1+n_2, p_c, p_t, T)$, then there is an $(n, poly(n)\cdot p_c, poly(n)\cdot p_t, \infty)$-algorithm $\mathcal{B}$ for any graph with n nodes with preprocessing time $O(t_p(O(m_0), p_c, p_t, T) \log n)$ where m_0 is a number of initial edges and update time $O(t_u(O(n), p_c, p_t, T) \log n + t_p(O(n), p_c, p_t, T)/T + \log^2 n)$. Moreover, if $\mathcal{A}$ is stable, then $\mathcal{B}$ has update time $O(t_u(O(n), p_c, p_t, T)+t_p(O(n), p_c, p_t, T)/T+ \log^2 n)$.*

Proof. By Lemma 4.23, we get a stable algorithm. That is, there is a stable $(n, O(Tp_c), p_t, T)$-algorithm $\mathcal{A}_0$ for a 3-bounded-degree graph where the number of edges is always between $n/2$ and $2n$ with preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(O(n), p_c, p_t, T) \log n + \log^2 n)$.

Let $U = n^2$. If $T \leq U$, then by Lemma 4.24, we get an algorithm that can handle long sequences of updates. That is, there is a stable $(n, \text{poly}(nU) \cdot p_c, \text{poly}(nU) \cdot p_t, U)$-algorithm $\mathcal{A}_1$ for a 3-bounded-degree graph where the number of edges is always between $n/2$ and $2n$ with preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) \log n+t_p(n, p_c, p_t, T)/T+\log^2 n)$. If $T > U$, then we can just treat $\mathcal{A}$ as $\mathcal{A}_1$. In either cases, we obtain $\mathcal{A}_1$.

By Lemma 4.25, we remove the condition that the graph has maximum degree 3. That is, there is a stable $(n, \text{poly}(nU)\cdot p_c, \text{poly}(nU)\cdot p_t, U)$-algorithm $\mathcal{A}_2$ for any graph with where the number of edges is always between $n/2$ and $2n$ with preprocessing time $O(t_p(n, p_c, p_t, T) + n) = O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) \log n + t_p(n, p_c, p_t, T)/T + \log^2 n)$.

Next, by Lemma 4.26, we remove the condition that the number of edges is not too small. That is, there is a stable $(n, \mathrm{poly}(nU)\cdot p_c, \mathrm{poly}(nU)\cdot p_t, U)$-algorithm $\mathcal{A}_3$ for any graph with n nodes where the number of edges is always at most $2n$, and $\mathcal{A}_3$ has preprocessing time $O(t_p(O(m_0), p_c, p_t, U))$ where m_0 is a number of initial edges, and update time $O(t_u(O(m), p_c, p_t, T)\log m + t_p(O(m), p_c, p_t, T)/T + \log^2 m)$ where m is a number of edges when update.

Next, by Lemma 4.27, we remove the condition that a graph is sparse. We obtain a stable $(n, \mathrm{poly}(nU)\cdot p_c, \mathrm{poly}(nU)\cdot p_t, U)$-algorithm $\mathcal{A}_4$ for any graph with n nodes with preprocessing time $O(t_p(O(m_0), p_c, p_t, T)\log n)$ where m_0 is a number of initial edges and update time $O(t_u(O(n), p_c, p_t, T)\log n + t_p(O(n), p_c, p_t, T)/T + \log^2 n)$.

Finally, by Lemma 4.28, we can handle update sequences of any length. That is, as $U = n^2$, we obtain an $(n, \mathrm{poly}(nU)\cdot p_c, \mathrm{poly}(nU)\cdot p_t, \infty)$-algorithm $\mathcal{A}_5$ for any graph with n nodes with preprocessing time $O(t_p(O(m_0), p_c, p_t, T)\log n)$ where m_0 is a number of initial edges and update time $O(t_u(O(n), p_c, p_t, T)\log n + t_p(O(n), p_c, p_t, T)/T + \log^2 n)$.

Note that if $\mathcal{A}$ is stable, then we do not need to use Lemma 4.23 in the first step, and we have that $\mathcal{A}_5$ has update time $O(t_u(O(n), p_c, p_t, T) + t_p(O(n), p_c, p_t, T)/T + \log^2 n)$. □

5 Global Expansion Decomposition

The backbone of both of our Monte Carlo and Las Vegas dynamic SF algorithms is the algorithm for finding *expansion decomposition* of a graph, which decomposes a graph into components of expanders and a sparse remaining part:

Theorem 5.1 (Global Expansion Decomposition Algorithm)**.** *There is a randomized algorithm $\mathcal{A}$ takes as inputs an undirected graph $G = (V, E)$ with $n \geq 2$ vertices and m edges and an expansion parameter $\alpha > 0$, and a failure probability parameter p. Then, in $O(m\gamma \log \frac{1}{p})$ time, $\mathcal{A}$ outputs two graphs $G^s = (V, E^s)$ and $G^d = (V, E^d)$ with the following properties:*

- *$\{E^s, E^d\}$ is a partition of E,*
- *G^s is sparse: $|E^s| \leq \alpha n\gamma$, and*
- *with probability $1 - p$, each connected component C of G^d either is a singleton or has high expansion, i.e. $\phi(C) \geq \alpha$.*

where the factor $\gamma = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$.

Throughout this section, we call the output of the algorithm above an *expansion decomposition*. There are two main steps for proving this theorem. First, we devise a near-linear time approximation algorithm for a problem called *most balanced sparse cut*. Second, we use our most balanced sparse cut algorithm as a main procedure for constructing the decomposition algorithm for Theorem 5.1.

Let us explain a high-level idea. Most balanced sparse cut problem is to find a cut S with largest number of nodes while $|S| \leq |V - S|$ such that $\phi(S) < \alpha$ where α is an input. This problem is closely related to sparsest cut and balanced cut problems (cf. [27]). One way to approximately solve both problems is by using the *cut-matching game* framework of Khandekar, Rao and Vazirani [26] together with known exact algorithms for the maximum flow problem. We modify this framework (in a rather straightforward way) so that (i) it gives a solution to the most balanced sparse cut problem and (ii) we can use *approximate* maximum flow algorithms instead of the exact ones. By plugging in near-linear time max flow algorithms [33, 35, 25], our algorithm runs in near-linear time. See Section 5.1 for details.

Then we use an approximate most balanced sparse cut algorithm as a main procedure for finding expansion decomposition. The main idea is simply to find the cut and recurse on both sides. However, the exact way for recursing is quite involved because we only have an approximate guarantee about the cut. Because of this, there is a factor of $\gamma = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$ in Theorem 5.1. See Section 5.2 for details.

5.1 Most-Balanced Sparse Cut

For any unweighted undirected graph $G = (V, E)$ with n vertices and m edges, we say that a cut $S \subset V$ is α-*sparse* if $\phi_G(S) < \alpha$. Let $\mathsf{OPT}(G, \alpha)$ be the *number of vertices* of the largest α-sparse cut S^* where $|S^*| \leq |V - S^*|$. If $\phi(G) \geq \alpha$, then we define $\mathsf{OPT}(G, \alpha) = 0$. Note that, if $\alpha \leq \alpha'$, then $\mathsf{OPT}(G, \alpha) \leq \mathsf{OPT}(G, \alpha')$. Now, we define most balanced sparse cut formally:

Definition 5.2 (Most Balanced Sparse Cut)**.** *For any unweighted undirected graph $G = (V, E)$ and parameters c_{size}, c_{exp} and α, a cut $S \subset V$ where $|S| \leq n/2$ is a (c_{size}, c_{exp})-approximate most-balanced α-sparse cut if $\phi_G(S) < \alpha$ and $|S| \geq \mathsf{OPT}(G, \alpha/c_{exp})/c_{size}$. That is, the size of S is as large as any (α/c_{exp})-sparse cut up to the c_{size} factor.*

Definition 5.3 (Most-Balanced Sparse Cut Algorithm)**.** *A (c_{size}, c_{exp})-approximate algorithm $\mathcal{A}$ for most-balanced sparse cut problem is given an unweighted undirected graph $G = (V, E)$ and a parameter α, and then $\mathcal{A}$ either*

- *finds a (c_{size}, c_{exp})-approximate most-balanced α-sparse cut S, or*
- *reports that $\phi(G) \geq \alpha/c_{exp}$.*

The main theorem in this section is the following.

Theorem 5.4. *There is a (c_{size}, c_{exp})-approximate most-balanced sparse cut algorithm with running time $\tilde{O}(m)$ where $c_{size} = \Theta(\log^2 n)$ and $c_{exp} = \Theta(\log^3 n)$.*

As the proof is by a straightforward modification of the cut-matching game framework by Khandekar, Rao and Vazirani [26] for solving balanced cut problem, we defer the proof to Appendix B.4.

5.2 The Global Decomposition Algorithm

In this section, we show how to construct the global expansion decomposition. Our fast construction is motivated by two algorithms which are too slow. However, they give main ideas of our algorithm so we discuss them below.

Early Attempts First, the following is a straight-forward recursive algorithm, which is too slow, for constructing the decomposition: use a sparsest cut algorithm to find an α-sparse cut S. If there is no α-sparse cut, then return the current graph as a component in G^d, otherwise include the cut edges $\partial(S)$ into G^s. Then recurse on $G[S]$ and $G[V \setminus S]$. In [3], they implicitly find a global expansion decomposition with this approach, since they focus on the space complexity and not running time. This algorithm can take $\Omega(mn)$ time because the cut S might be very small, so the recursion tree can be very unbalanced.

Second, a next natural approach is to use a most-balanced sparse cut (recall Definition 5.2) instead of a sparsest cut. Suppose that there exists a $(2, 1)$-approximate most-balanced sparse cut algorithm, this approach would have been straightforward:

Proposition 5.5. *If there is a $(2,1)$-approximate most-balanced sparse cut algorithm $\mathcal{A}$ with running time $T(m)$ on a graph with n vertices and m edges, then there is an algorithm for computing expansion decomposition for graph with n vertices and m edges in $O(T(m,n)\log n)$ time.*

Proof. For any graph $H = (V_H, E_H)$, the recursive procedure $\mathsf{decomp}(H)$ does the following: if $\mathcal{A}(H,\alpha)$ reports that $\phi(H) \geq \alpha$ then return H as a connected components in G^d. Else $\mathcal{A}(H,\alpha)$ finds a cut S where $\mathsf{OPT}(H,\alpha)/2 \leq |S| \leq \mathsf{OPT}(H,\alpha)$. Add the cut edges $\partial_H(S)$ into G^s, then recurse on $\mathsf{decomp}(H[S])$ and $\mathsf{decomp}(H[V_H \setminus S])$. Given an input graph G, we claim that $\mathsf{decomp}(G)$ returns an expansion decomposition of G.

To see the correctness, 1) each component H in G^d has high expansion: $\phi(H) \geq \alpha$ by construction, and 2) the number of edges in G^s is at most $O(\alpha n \log n)$ because whenever $\partial_H(S)$ is added into G^s, we charge these edges to nodes in S. By averaging, each node is charged $\delta_H(S)/|S| = \phi_H(S) \leq \alpha$ Since $|S| \leq |V_H|/2$, each node can be charged $O(\log n)$ times.

To bound the running time, it is enough to show that the recursion depth is $O(\log n)$ because the graphs in the two branches of the recursion $H[S]$ and $H[V_H \setminus S]$ are disjoint. Let S' be such that $|S'| = \mathsf{OPT}(H[V_H \setminus S], \alpha)$. The case where $|S| \geq |V_H|/4$ or $|S'| \geq |V_H|/4$ can happen at most $O(\log n)$ times throughout the recursion. Now, suppose that $|S| < |V_H|/4$ and $|S'| < |V_H|/4$. So $|S \cup S'| < |V_H|/2$. We claim that $|S'| \leq \mathsf{OPT}(H,\alpha)/2$ which can also happen at most $O(\log n)$ times. Suppose otherwise, then $S \cup S'$ is an α-sparse cut in H where $|V_H|/2 \geq |S \cup S'| > \mathsf{OPT}(H,\alpha)/2 + \mathsf{OPT}(H,\alpha)/2 = \mathsf{OPT}(H,\alpha)$ which is a contradiction. So the recursion depth is $O(\log n)$. ☐

Unfortunately, we do not have an efficient $(2,1)$-approximate most-balanced sparse cut algorithm. By Theorem 5.4, we only have $(O(\log^2 n), O(\log^3 n))$-approximate most-balanced sparse cut algorithm. Our main technical goal is, then, to use this weaker guarantee while making this approach goes through.

Efficient Algorithm Let $\mathcal{A}'$ be the algorithm from Theorem 5.4. Recall that, given any graph H with n' nodes and a parameter α', $\mathcal{A}'$ has approximation ratio $c_{size}(n') = \Theta(\log^2 n')$ and $c_{exp}(n') = \Theta(\log^3 n')$, i.e., with probability $1 - \frac{1}{(n')^k}$ for a large constant k, $\mathcal{A}'$ either outputs an α'-sparse cut S where $\mathsf{OPT}(H, \alpha'/c_{exp}(n'))/c_{size}(n') \leq |S| \leq n'/2$ or reports that $\phi(H) \geq \alpha'/c_{exp}(n')$. In the following, sometimes the input graph H is small and so the probability $1 - \frac{1}{(n')^k}$ is high enough for our purpose. To boost this probability to $1 - p'$ for arbitrarily small p, we define the algorithm $\mathcal{A}$ to be the algorithm that repeatedly runs $\mathcal{A}'$ for $O(\log \frac{1}{p'})$ iterations. If there is some iteration where $\mathcal{A}'$ returns an α'-sparse cut S, then $\mathcal{A}$ just return S. Otherwise, there is no iteration that $\mathcal{A}'$ returns an α'-sparse cut, then $\mathcal{A}$ report that $\phi(H) \geq \alpha'/c_{exp}(n')$ with probability $1 - p'$. In the following, we will use this modified algorithm $\mathcal{A}$ instead of the algorithm $\mathcal{A}'$ from Theorem 5.4.

Fact 5.6. *For any given p, the algorithm $\mathcal{A}$ defined above works correctly with probability $1 - p$.*

Suppose that the input of our decomposition is a graph G with n vertices and m edges, a parameter α, and the target success probability $1 - p$. We now define some parameters. Let $\bar{c}_{size} = c_{size}(n)$, $\bar{c}_{exp} = c_{exp}(n)$, and $\epsilon = \sqrt{\log \bar{c}_{exp}/\log n}$. Let $\bar{s}_1, \ldots, \bar{s}_L$ be such that $\bar{s}_1 = n/2 + 1$, $\bar{s}_L \leq 1$, and $\bar{s}_\ell = \bar{s}_{\ell-1}/n^\epsilon$ for $1 < \ell < L$. Hence, $L \leq 1/\epsilon$. Let $\alpha_1, \ldots, \alpha_L$ be such that $\alpha_L = \alpha \bar{c}_{exp}$ and $\alpha_\ell = \alpha_{\ell+1}\bar{c}_{exp}$ for $\ell < L$. Hence, $\alpha_1 = \alpha \bar{c}_{exp}^L$. Let $p' = p/n^2$. So $\mathcal{A}$ repeatedly runs $\mathcal{A}'$ for $O(\log \frac{1}{p'}) = O(\log \frac{n}{p})$ iterations.

For any graphs $H = (V_H, E_H)$, $I = (V_I, E_I)$ and a number ℓ, the main procedure $\mathsf{decomp}(H, I, \ell)$ is defined as in Algorithm 5.1.

We claim that $\mathsf{decomp}(G, G, 1)$ returns an expansion decomposition for G with parameter α in time $\tilde{O}(m^{1+o(1)})$. The following lemmas show the correctness. For readability, we first assume that $\mathcal{A}$ works correctly with certainty, and then we remove the assumption later.

Given that, $\mathsf{OPT}(I, \alpha_\ell) < \bar{s}_\ell$ do the following:

1. If H is a singleton or $\mathcal{A}(H, \alpha_\ell)$ reports $\phi(H) \geq \alpha_\ell / c_{exp}(|V_H|)$, then return H as a connected components in G^d.
2. Else, $\mathcal{A}(H, \alpha_\ell)$ returns an α_ℓ-sparse cut S in H where $\mathsf{OPT}(H, \alpha_\ell / c_{exp}(|V_H|))/c_{size}(|V_H|) \leq |S| \leq |V_H|/2$.
   (a) If $|S| \geq \bar{s}_{\ell+1}/\bar{c}_{size}$, then add the cut edges $\partial_H(S)$ into G^s and recurse on $\mathsf{decomp}(H[S], H[S], 1)$ and $\mathsf{decomp}(H[V_H \setminus S], I, \ell)$.
   (b) Else, recurse on $\mathsf{decomp}(H, H, \ell + 1)$.

Algorithm 5.1: $\mathsf{decomp}(H, I, \ell)$ where $H = (V_H, E_H)$ and $I = (V_I, E_I)$.

Proposition 5.7. *For any graphs $H, I \subseteq G$ and ℓ, if $\mathsf{decomp}(H, I, \ell)$ is called, then the invariant $\mathsf{OPT}(I, \alpha_\ell) < \bar{s}_\ell$ is satisfied.*

Proof. When $\ell = 1$, $\mathsf{OPT}(I, \alpha_1) < \bar{s}_1 = n/2+1$ for any I in a trivial way. In particular, the invariant is satisfied when $\mathsf{decomp}(G, G, 1)$ or $\mathsf{decomp}(H[S], H[S], 1)$ is called. The invariant $\mathsf{decomp}(H[V_H \setminus S], I, \ell)$ is the same as the one for $\mathsf{decomp}(H, I, \ell)$, and hence is satisfied. Finally, $\mathsf{decomp}(H, H, \ell+1)$ is called in step 2.b only when $|S| < \bar{s}_{\ell+1}/\bar{c}_{size}$, so we have

$$\begin{aligned} \mathsf{OPT}(H, \alpha_{\ell+1}) &= \mathsf{OPT}(H, \alpha_\ell/\bar{c}_{exp}) \leq \mathsf{OPT}(H, \alpha_\ell/c_{exp}(|V_H|)) \\ &\leq c_{size}(|V_H|)|S| < \frac{c_{size}(|V_H|)}{\bar{c}_{size}} \leq \bar{s}_{\ell+1}. \end{aligned}$$

These inequalities hold because $c_{exp}(|V_H|) \leq \bar{c}_{exp}$ and $c_{size}(|V_H|) \leq \bar{c}_{size}$ as $|V_H| \leq n$, and we know that $\mathsf{OPT}(H, \alpha) \leq \mathsf{OPT}(H, \alpha')$ for any $\alpha \leq \alpha'$. □

Proposition 5.8. *At level L, $\mathsf{decomp}(H, H, L)$ always return the graph H (no further recursion).*

Proof. The invariant $\mathsf{OPT}(H, \alpha_L) < \bar{s}_L \leq 1$ implies that $\phi(H) > \alpha_L$. Hence, $\mathcal{A}(H, \alpha_L)$ can never find an α_L-sparse cut. So $\mathcal{A}(H, \alpha_L)$ must report $\phi(H) \geq \alpha_L / c_{exp}(|V_H|)$. So H is returned. □

Proposition 5.9. *Each component H in G^d either is a singleton or has expansion $\phi(H) \geq \alpha$.*

Proof. H is returned as a component in G^d only when H is a singleton or $\mathcal{A}$ reports that $\phi(H) \geq \alpha_\ell / c_{exp}(|V_H|) \geq \alpha_L/\bar{c}_{exp} = \alpha$. Note that $\ell < L$ by Proposition 5.8. □

Proposition 5.10. *The number of edges in G^s is at most $O(\alpha_1 n \log n)$.*

Proof. Suppose $\partial_H(S)$ is added into G^s when $\mathsf{decomp}(H, I, \ell)$ for some graphs $H, I \subset G$, we charge these edges to nodes in S. By averaging, each node is charged $\delta_H(S)/|S| = \phi_H(S) \leq \alpha_\ell \leq \alpha_1$. Since $|S| \leq |V_H|/2$, each node can be charged $O(\log n)$ times. □

To analyze the running time, we need some more notation. We define the recursion tree $\mathcal{T}$ of $\mathsf{decomp}(G, G, 1)$ as follows. Each node of $\mathcal{T}$ represents the parameters of the procedure $\mathsf{decomp}(H, I, \ell)$. That is, $(G, G, 1)$ is the root node. For each (H, I, ℓ), if $\mathsf{decomp}(H, I, \ell)$ returns H, then (H, I, ℓ) is a leaf. If $\mathsf{decomp}(H, I, \ell)$ recurses on $\mathsf{decomp}(H[S], H[S], 1)$ and $\mathsf{decomp}(H[V \setminus S], I, \ell)$, then $(H[S], H[S], 1)$ and $(H[V \setminus S], I, \ell)$ is a left and right children of (H, I, ℓ) respectively, and the edges $((H, I, \ell), (H[S], H[S], 1))$ and $((H, I, \ell), (H[V \setminus S], I, \ell))$ are *left edge* and *right edge*

respectively. If $\mathsf{decomp}(H, I, \ell)$ recurses on $\mathsf{decomp}(H, H, \ell+1)$, then $(H, H, \ell+1)$ is the only child of (H, I, ℓ), and the edge $((H, I, \ell), (H, H, \ell+1))$ is a *down edge*.

Lemma 5.11. *If the depth of the recursion tree $\mathcal{T}$ is D, then the total running time is $\tilde{O}(mD)$.*

Proof. The running time on $\mathsf{decomp}(H, I, \ell)$, excluding the time spent in the next recursion level, is at most $\tilde{O}(|E_H| \log \frac{1}{p})$ because the algorithm $\mathcal{A}$ just repeatedly runs the algorithm $\mathcal{A}'$ from Theorem 5.4 for $O(\log \frac{1}{p})$ times, and $\mathcal{A}'$ runs in near-linear time.

Let $\mathcal{T}_d$ be the set of all the nodes in $\mathcal{T}$ of depth *exactly* d. We can see that, for any two nodes $(H, I, \ell), (H', I, \ell') \in \mathcal{T}_d$, H and H' are disjoint subgraphs of G. So, the total running time we spent on $\mathcal{T}_d$ is $\tilde{O}(m \log \frac{1}{p'})$. Hence, the total running time of $\mathcal{T}$ is $\tilde{O}(mD \log \frac{1}{p'})$. □

Lemma 5.12. *Let P be any path from leaf to root of $\mathcal{T}$. P contains at most $\log n$ left edges, $L \log n$ down edges, and $L \log n \times \gamma$ right edges where $\gamma = \bar{c}_{size} n^{\epsilon}$. Therefore, $\mathcal{T}$ has depth at most $L \log n \cdot \bar{c}_{size} n^{\epsilon}$.*

Proof. Consider the left edge $((H, I, \ell), (H[S], H[S], 1))$. Since $|S| \le |V_H|/2$, P contains $\log n$ left edges. Between any two left edges in P there are at most L down edges be Proposition 5.8. So P contains at most $L \log n$ down edges. To prove that there are at most $L \log n \times \gamma$ right edges in P, it suffices to prove that there cannot be γ right edges between any left edges or down edges in P.

Suppose that $(H_1, I, \ell), \ldots, (H_\gamma, I, \ell)$ are nodes in P where, for each i, $((H_i, I, \ell), (H_{i+1}, I, \ell))$ is a right edge, and (H_1, I, ℓ) is a deeper endpoint of left edges or a down edges (hence $I = H_1$).

For each i, let S_i be the cut such that $H_{i+1} = H_i[V_{H_i} \setminus S_i]$ and $\phi_{H_i}(S_i) < \alpha_\ell$. Since $\{S_i\}_i$ are mutually disjoint and $\partial_{H_1}(\bigcup_{i=1}^{\gamma} S_i) \subset \bigcup_{i=1}^{\gamma} \partial_{H_i}(S_i)$, we can conclude $\phi_{H_1}(\bigcup_{i=1}^{\gamma} S_i) < \alpha_\ell$. However, we also have that $|S_i| \ge \bar{s}_{\ell+1}/\bar{c}_{size}$, for all i, and hence $|\bigcup_{i=1}^{\gamma} S_i| \ge \gamma \bar{s}_{\ell+1}/\bar{c}_{size} \ge n^{\epsilon} \bar{s}_{\ell+1} = \bar{s}_\ell$. So $\bigcup_{i=1}^{\gamma} S_i$ contradicts the invariant for $\mathsf{decomp}(H_1, H_1, \ell)$ (note that $I = H_1$) which says $\mathsf{OPT}(H_1, \alpha_\ell) < \bar{s}_\ell$. □

Now, we can conclude the theorem.

Proof of Theorem 5.1. Recall that $\epsilon = \sqrt{\log \bar{c}_{exp}/\log n} = \Theta(\sqrt{\log\log n/\log n})$ and $L \le 1/\epsilon$. Given G and α, by Lemma 5.11 and Lemma 5.12 the running time of $\mathsf{decomp}(G, G, 1)$ is $\tilde{O}(mL\bar{c}_{size} n^{\epsilon} \log \frac{1}{p'}) = O(m^{1+O(\sqrt{\log\log n/\log n})} \log \frac{1}{p})$. By Proposition 5.9 and Proposition 5.10, for each connected component H of G^d, H is a singleton or $\phi(H) \ge \alpha$ and the number of edges in G^s is $O(\alpha_1 n \log n) = O(\alpha c_{exp}^L n \log n) = \alpha n^{1+O(\sqrt{\log\log n/\log n})}$. This is the output for the expansion decomposition as claimed.

Finally, we remove the assumption that $\mathcal{A}$ is deterministic, and show that with probability $1 - p$, the outputted decomposition is correct. Observe that, as the depth of the recursion tree $\mathcal{T}$ of $\mathsf{decomp}(G, G, 1)$ is $D = L \log n \cdot \bar{c}_{size} n^{\epsilon}$, then $\mathcal{A}$ is called at most nD times. So the probability that $\mathcal{A}$ always works correctly is at least $1 - p'nD \ge 1 - p$. □

Finally, we add some easy observation that will be used in Section 6.

Proposition 5.13. *All the edges added into E^s by Algorithm 5.1 are the cut edges $\partial_H(S)$ for some subgraph H of G and a cut $S \subset V(H)$ where $\phi_H(S) \le \alpha_1 = \alpha \bar{c}_{exp}^L = \alpha\gamma$ where $\gamma = n^{O(\sqrt{\log\log n/\log n})}$.*

6 Local Expansion Decomposition

In this section, we show a novel "local" version of the global expansion decomposition algorithm from Section 5. This algorithm differs from the algorithm from Section 5 in two aspects. First, it needs that the input graph G is obtained by deleting some edge set D from another graph G_b, which "we expect that" $\phi(G_b) \geq \alpha_b$. Second, the algorithm is *local,* in the sense that its running time essentially depends only on the size of $|D|$ and α_b, and not the size of G.

This algorithm is a crucial tool for obtaining our Las Vegas dynamic SF algorithm. It is important that the dependency of $|A|$ in the running time is truly subquadratic (i.e. $|A|^{2-\epsilon}$ for some ϵ), which is the case here. For any n, recall $\gamma = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$ is the factor from Theorem 5.1.

Theorem 6.1 (Local Expansion Decomposition Algorithm)**.** *For any constant $\epsilon \in (0,1)$, there is an algorithm $\mathcal{A}$ that can do the following:*

- *$\mathcal{A}$ is given pointers to $n, p, \Delta, G_b, D, \alpha_b$ which are stored in a memory: p is a failure probability parameter. $G_b = (V, E_b)$ is a Δ-bounded degree graphs with n nodes (represented by an adjacency list). α_b is an expansion parameters where $\alpha_b < \frac{1}{\gamma^{\omega(1)}}$. $D \subset E_b$ is a set of edges in G_b. Let A denote a set of endpoints of edges in D. Let $G = (V, E) = (V, E_b - D)$ be the graph that $\mathcal{A}$ will compute the decomposition on.*
- *Then (without reading the whole inputs), in time $\tilde{O}(\frac{\Delta^{5.5}|A|^{1.5+\epsilon}}{\alpha_b^{4+\epsilon}\epsilon}\gamma \log \frac{1}{p})$, $\mathcal{A}$ either*
  - *reports that G_b has expansion less than α_b, or*
  - *outputs 1) a set of edges $E^s \subset E$, and 2) all components, except the largest one, of G^d where $G^d = (V, E^d) = (V, E - E^s)$.*
- *Moreover, if G_b has expansion at least α_b, then with probability $1 - O(p)$ we have*
  - *$|E^s| \leq 4\Delta|A|/\alpha_b$, and*
  - *each connected component C of G^d either is a singleton or has high expansion: $\phi(C) \geq \alpha$ where $\alpha = (\alpha_b/6\Delta)^{1/\epsilon}$.*

Observe that ϵ is a trade-off parameter such that, on one hand when ϵ is small, the algorithm is fast but has a bad expansion guarantee in the output, on the other hand when ϵ is big, the algorithm is slow but has a good expansion guarantee.

Remark 6.1. In Theorem 6.1, by *outputting a components C* of G^d, it means for each nodes $u \in V(C)$, (u, C) is outputted indicating that $u \in V(C)$. Since the algorithm $\mathcal{A}$ in Theorem 6.1 outputs all components, except one, of G^d. We can infer the size of *all* components in G^d, and for any node $u \in V$, sa component C of G^d where $u \in V(C)$.

Although the main idea for proving Theorem 6.1 is similar to Theorem 5.1, the analysis for is more involved. We explain the high level idea below. There are two main steps. First, in Section 6.1 we show an algorithm for a problem called *most-balanced sparse cut with respect to overlapping cuts (MSO cut)* , which is basically a "local version" of most-balanced sparse cut defined in Section 5.1. Second, we use this algorithm as a main procedure for constructing the decomposition in Section 6.2.

In MSO cut problem, we are given a graph $G = (V, E)$, a target set $A \subset V$ and a parameter α. Then we need to find a α-sparse cut S (i.e. $\phi_G(S) < \alpha$) where $|S| \leq |V - S|$ such that $|S|$ is larger than all α-sparse cuts which are "near" the target set A. ("Nearness" is defined precisely in

Definition 6.2). To compare, in most-balanced sparse cut problem, $|S|$ needs to be larger than *all* α-sparse cuts. The nice thing about MSO cut problem is that there is an approximate algorithm for this problem which is "local" (i.e. its running time depends essentially only on $vol(A)$, and not $|V|$). This algorithm can be obtained quite easily by slightly modifying and analyzing the algorithm by Orecchia and Zhu [30] for local cut improvement problem. See Section 6.1 for details.

Next, to obtain the decomposition, the main approach is to find an approximate MSO cut where the target set A is the endpoints of D (union with some additional nodes) and recurse on both sides. However, the analysis is more involved than the one in Theorem 5.1. One important observation is that, when $\phi(G_b) \geq \alpha_b$, any $(\frac{\alpha_b}{2})$-sparse cut S in $G = G_b - D$ must be "near" to D. Intuitively, this is because the expansion of S get halved after deleting edges in D. This justifies why it is enough to find approximate $(\frac{\alpha_b}{2})$-sparse MSO cuts instead of finding approximate most-balanced sparse cuts as in the algorithm for Theorem 5.1. As our approximate MSO cut algorithm is local, we can output the decomposition in time essentially independent from the size of G. See Section 6.2 for details.

6.1 Most-balanced Sparse Cut with respect to Overlapping Cuts

In this section, we show the crucial tool for proving Theorem 6.1. First, we need this definition:

Definition 6.2 (Overlapping)**.** *For any graph $G = (V, E)$ and set $A \subset V$, a cut $S \subset V$ is (A, σ)-overlapping in G if $|S \cap A|/|S| \geq \sigma$.*

Let $G = (V, E)$ be a graph. Recall that a cut S is α-sparse if $\phi(S) = \frac{\delta(S)}{\min\{|S|,|V-S|\}} < \alpha$. For any set $A \subset V$, an overlapping parameter σ and an expansion parameter α, let S^* be the largest α-sparse (A, σ)-overlapping cut where $|S^*| \leq |V - S^*|$. We define $\mathsf{OPT}(G, \alpha, A, \sigma) = |S^*|$. If S^* does not exist, then $\mathsf{OPT}(G, \alpha, A, \sigma) = 0$. Now, we define most balanced sparse cut problem formally:

Definition 6.3 (Most-Balanced Sparse Cut w.r.t Overlapping Cuts)**.** *For any graph $G = (V, E)$, a set $A \subset V$, and parameters $c_{size}, c_{exp}, \sigma$ and α, a cut S where $|S| < |V - S|$ is a (c_{size}, c_{exp})-approximate most-balanced α-sparse cut w.r.t. (A, σ)-overlapping cuts if $\phi'(S) < \alpha$ and $|S| \geq \mathsf{OPT}(G, \alpha/c_{exp}, A, \sigma)/c_{size}$. We also write S is a $(c_{size}, c_{exp}, \alpha, A, \sigma)$-MSO cut.*

We note that $(c_{size}, c_{exp}, \alpha, A, \sigma)$-MSO cut S may not be (A, σ)-overlapping. The existence of S just show that for any (A, σ)-overlapping cut of size at least $c_{size}|S|$ must have expansion at least α/c_{exp}.

Definition 6.4 (Most-Balanced Sparse Cut w.r.t Overlapping Cuts Algorithm)**.** *A (c_{size}, c_{exp})-approximate algorithm $\mathcal{A}$ for most-balanced sparse cut w.r.t. overlapping cuts problem is given a graph $G = (V, E)$, a set $A \subset V$, and an overlapping parameter σ and an expansion parameter α , and then $\mathcal{A}$ either*

- *finds a $(c_{size}, c_{exp}, \alpha, A, \sigma)$-MSO cut S, or*
- *reports that there is no (α/c_{exp})-sparse (A, σ)-overlapping cut.*

We also write that $\mathcal{A}$ is a (c_{size}, c_{exp})-approximate MSO cut algorithm.

The main result of this section is the following:

Theorem 6.5. *Given a graph $G = (V, E)$, a set $A \subset V$ where $|A| \leq 4|V - A|$, an expansion parameter α, and an overlapping parameter $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$, there is a (c_{size}, c_{exp})-approximate MSO cut algorithm with running time $O((vol(A)/\sigma)^{1.5} \log^2(\frac{vol(A)}{\sigma}))$ where $c_{size} = 3/\sigma$ and $c_{exp} = 3/\sigma$.*

This algorithm can be obtained quite easily by slightly modifying and analyzing the algorithm by Orecchia and Zhu [30] for local cut improvement problem. Therefore, we defer the proof to Appendix B.5.

6.2 The Local Decomposition Algorithm

Throughout this section, let n and p be given. Let $G_b = (V, E_b)$ be a graph with n nodes with maximum degree Δ. Let $D \subset E$ be some set of edges. Let $G = (V, E) = (V, E_b - D)$ and A be a set of endpoints of edges in D. We call G_b a *before graph*. Our algorithm is given D, a constant ϵ, and an expansion parameter α_b as inputs. We want the algorithm to output a set of edges $E^s \subset E$ and all components, except the largest one, of G^d where $G^d = (V, E^d) = (V, E - E^s)$. Let $G^s = (V, E^s)$. Let $\bar{t} = \tilde{O}(\frac{\Delta^{5.5}|A|^{1.5+\epsilon}}{\alpha_b^{4+\epsilon}\epsilon}\gamma \log \frac{1}{p})$ be the parameter indicating the time limit.

We now define some more notations. Let $\mathcal{A}_{cut}$ be the deterministic algorithm for finding MSO cuts from Theorem 6.5. We set the overlapping parameter $\sigma = \alpha_b/2\Delta$ for $\mathcal{A}_{cut}$. Hence the approximation ratios of $\mathcal{A}_{cut}$ are $c_{size} = 3/\sigma$ and $c_{exp} = 3/\sigma$. Let $\mathcal{A}_{decomp}$ be the randomized algorithm for finding an expansion decomposition from Theorem 5.1. We set $p' = p/n^2$ as the failure probability parameter for $\mathcal{A}_{decomp}$.

The following notations was similarly defined as in Section 5.2. Let $\bar{s}_1, \ldots, \bar{s}_L$ be such that $\bar{s}_1 = 4|A|/\alpha_b + 1$, $\bar{s}_L \leq 1$, and $\bar{s}_\ell = \bar{s}_{\ell-1}/(\bar{s}_1)^\epsilon$ for $1 < \ell < L$. Hence, $L \leq 1/\epsilon$. We denote $\alpha = (\alpha_b/6\Delta)^{1/\epsilon}$. Let $\alpha_1, \ldots, \alpha_L$ be such that $\alpha_L = \alpha$ and $\alpha_\ell = \alpha_{\ell+1}c_{exp}$ for $\ell < L$. Hence, $\alpha_1 = \alpha c_{exp}^{L-1}$.

Fact 6.6. $\alpha_1 < \alpha_b/2$.

Proof. We have

$$\alpha_1 = \alpha c_{exp}^{L-1} = (\frac{\alpha_b}{6\Delta})^{1/\epsilon}(\frac{3}{\sigma})^{1/\epsilon-1} \leq (\frac{\alpha_b}{6\Delta})^{1/\epsilon}(\frac{6\Delta}{\alpha_b})^{1/\epsilon-1} = \alpha_b/6\Delta < \alpha_b/2.$$

□

For any graphs $H = (V_H, E_H)$, $I = (V_I, E_I)$, and a number ℓ, the main procedure $\mathsf{decomp}(H, I, \ell)$ is defined as in Algorithm 6.1. For any α' and $B \subset V_H$, recall that $\mathsf{OPT}(H, \alpha')$ is the size of the largest α'-sparse cut S in H where $|S| \leq |V_H - S|$, and $\mathsf{OPT}(H, \alpha', B, \sigma)$ is the size of the largest α'-sparse (B, σ)-overlapping cut S in H where $|S| \leq |V_H - S|$. By definition, $\mathsf{OPT}(H, \alpha') \geq \mathsf{OPT}(H, \alpha', B, \sigma)$.

The algorithm is simply to run $\mathsf{decomp}(G, G, 1)$ with time limit $\bar{t}$. If $\mathsf{decomp}(G, G, 1)$ takes time more than $\bar{t}$, then we reports that $\phi(G_b) < \alpha_b$.

6.2.1 Validity

We first show that the parameters for $\mathcal{A}_{cut}$ are valid when it is called.

Lemma 6.7. *Whenever $\mathcal{A}_{cut}(H, \alpha_\ell, B_H, \sigma)$ is called, we have that $\sigma \geq 3\frac{|B_H|}{|V_H - B_H|}$ satisfying the requirement for $\mathcal{A}_{cut}$ as stated in Theorem 6.5.*

Proof. Observe that $\mathcal{A}_{cut}$ can be called only when the condition in Step 2 of Algorithm 6.1 is false: $|V_H - B_H| \geq \frac{3}{\sigma}|B_H|$. That is, $\sigma \geq 3\frac{|B_H|}{|V_H - B_H|}$. □

Next, we show that if $\mathsf{decomp}(G, G, 1)$ finishes with in time limit $\bar{t}$, then it outputs 1) a set of edges $E^s \subset E$, 2) all components, except the largest one, of G^d where $G^d = (V, E^d) = (V, E - E^s)$. Recall Remark 6.1, by outputting a components C of G^d, it means for each nodes $u \in V(C)$, (u, C) is outputted indicating that $u \in V(C)$.

It is clear from Algorithm 6.1 that when E^s is outputted, we have $G^d = (V, E^d) = (V, E - E^s)$. It is also obvious how to prevent Algorithm 6.1 from outputting the largest component C of G^d,

1. Set $B_H = (A \cup A_H) \cap V_H$ where A_H is the set of endpoints of edges in $\partial_G(V_H)$
2. If $|V_H - B_H| < \frac{3}{\sigma}|B_H|$, then run $\mathcal{A}_{decomp}(H, \alpha, p')$ which partitions E_H into E_H^s and E_H^d. Add E_H^s into G^s, and return each connected component of a graph induced by E_H^d as a components in G^d.
3. If $\ell = L$, then return H as a connected components in G^d.
4. If H is a singleton or $\mathcal{A}_{cut}(H, \alpha_\ell, B_H, \sigma)$ reports there is no (α_ℓ/c_{exp})-sparse (B_H, σ)-overlapping cut, then return H as a connected components in G^d.
5. Else, $\mathcal{A}_{cut}(H, \alpha_\ell, B_H, \sigma)$ returns an α_ℓ-sparse cut S in H where $\mathsf{OPT}(H, \alpha_\ell/c_{exp}, B_H, \sigma)/c_{size} \le |S| \le |V_H|/2$.
   (a) If $|S| \ge \bar{s}_{\ell+1}/c_{size}$, then add the cut edges $\partial_H(S)$ into G^s and recurse on $\mathsf{decomp}(H[S], H[S], 1)$ and $\mathsf{decomp}(H[V_H \setminus S], I, \ell)$.
   (b) Else, recurse on $\mathsf{decomp}(H, H, \ell + 1)$.

Algorithm 6.1: $\mathsf{decomp}(H, I, \ell)$ where $H = (V_H, E_H)$ and $I = (V_I, E_I)$

which may take a lot of time. That is, whenever Algorithm 6.1 return H as a component of G^d in Step 2,3 or 4, we do nothing if H is "reached" by recursing from only the "larger" side of the cut. Otherwise, we list all nodes $u \in V(H)$, and output (u, H).

6.2.2 Correctness

In this section, we prove that if $\phi(G_b) \ge \alpha_b$, then the outputs of $\mathsf{decomp}(G, G, 1)$ have the desired properties. For readability, we first assume that $\mathcal{A}_{decomp}$ works correctly with certainty, and then we remove the assumption later.

To bound $|E^s|$, we first need the following fact:

Proposition 6.8. *Suppose that $\phi(G_b) \ge \alpha_b$. For any $\alpha' < \alpha_b/2$, any α'-sparse cut S in G must have size at most $|S| \le 4|A|/\alpha_b$.*

Proof. Suppose otherwise that $|S| > 4|A|/\alpha_b$. Then, as $G = G_b - D$, we have

$$\delta_G(S) \ge \delta_{G_b}(S) - |D| \ge \alpha_b|S| - 2|A| > \alpha_b|S| - \alpha_b|S|/2 = \alpha_b|S|/2,$$

which means, $\phi_G(S) > \alpha_b/2 > \alpha'$, a contradiction. □

Lemma 6.9. *If $\phi(G_b) \ge \alpha_b$, then $|E^s| \le 4\Delta|A|/\alpha_b$.*

Proof. Observe all the edges added into E^s by Algorithm 6.1 are the cut edges $\partial_H(S)$ for some subgraph H of G and a cut $S \subset V(H)$. Let us list all sets of edges $\partial_{H_1}(S_1), \ldots, \partial_{H_t}(S_t)$ that constitute E^s.

Claim 6.10. For all i, $\phi_{H_i}(S_i) < \alpha_b/2$.

Proof. The cut edges $\partial_{H_i}(S_i)$ are added into E^s by Algorithm 6.1 either in Step 2 or Step 5.a. In the first case, we know from Proposition 5.13 that $\phi_{H_i}(S_i) < \alpha\gamma$. Now, as $\alpha = O(\alpha_b/\Delta)^{1/\epsilon}$ and ϵ is less than 1 by some constant, we have $\alpha \le \alpha_b^{1-\delta}$ for some constant $\delta > 0$. Since $\alpha_b \le \frac{1}{\gamma^{\omega(1)}}$, we have

$\gamma \leq \frac{1}{\alpha_b^{o(1)}}$. Therefore, $\phi_{H_i}(S_i) < \alpha\gamma \leq \alpha_b^{1-\delta+o(1)} < \alpha_b/2$. In the second case, we know from Step 5 that $\phi_{H_i}(S_i) < \alpha_1 < \alpha_b/2$. □

Observe that for any S_i and S_j where $i \neq j$, either $S_i \subset S_j$, $S_j \subset S_i$ or $S_i \cap S_j = \emptyset$. We say that S_i is *maximal* if there is no $j \neq i$ where $S_j \supset S_i$. Let $S'_1, \ldots, S'_{t'}$ be all the maximal sets. By the maximality, we have that $S'_1, \ldots, S'_{t'}$ are mutually disjoint, and $|E^s| \leq vol(\bigcup_{i=1}^{t'} S'_i)$. Let $H'_1, \ldots, H'_{t'}$ be the corresponding subgraphs of G where $\partial_{H'_{i'}}(S'_{i'}) \subset E^s$. By choosing the appropriate ordering, we have write $H'_1 = G$, $H'_2 = G[V - S'_1], \ldots, H'_{t'} = G[V - \bigcup_{i=1}^{t'-1} S'_i]$.

Claim 6.11. $\delta_G(\bigcup_{i=1}^{t'} S'_i) \leq \sum_{i=1}^{t'} \delta_{H'_i}(S'_i)$

Proof. We will prove that $\partial_G(\bigcup_{i=1}^{t'} S'_i) \subseteq \bigcup_{i=1}^{t'} \partial_{H'_i}(S'_i)$. Let $(u,v) \in \partial_G(\bigcup_{i=1}^{t'} S'_i)$. Suppose that $u \in S'_j$ for some j. Then $v \in V - \bigcup_{i=1}^{t'} S'_i \subseteq V(H'_j) - S'_j$ because $V(H'_j) = V - \bigcup_{i=1}^{j-1} S'_i$. Therefore, $(u,v) \in \partial_{H'_j}(S'_j) \subset \bigcup_{i=1}^{t'} \partial_{H'_i}(S'_i)$. □

By the two claims above, we have that $\phi_G(\bigcup_{i=1}^{t'} S'_i) < \alpha_b/2$. This is because

$$\begin{aligned}\phi_G(\bigcup_{i=1}^{t'} S'_i) &= \frac{\delta_G(\bigcup_{i=1}^{t'} S'_i)}{|\bigcup_{i=1}^{t'} S'_i|} = \frac{\delta_G(\bigcup_{i=1}^{t'} S'_i)}{\sum_{i=1}^{t'} |S'_i|} \leq \frac{\sum_{i=1}^{t'} \delta_{H'_i}(S'_i)}{\sum_{i=1}^{t'} |S'_i|} \\ &\leq \max_{i \leq t'} \frac{\delta_{H'_i}(S'_i)}{|S'_i|} = \max_{i \leq t'} \phi_{H'_i}(S'_i) < \alpha_b/2.\end{aligned}$$

Therefore, by Proposition 6.8 we have

Claim 6.12. $|\bigcup_{i=1}^{t'} S'_i| \leq 4|A|/\alpha_b$.

so $|E^s| \leq vol(\bigcup_{i=1}^{t'} S'_i) \leq \Delta|\bigcup_{i=1}^{t'} S'_i| \leq 4\Delta|A|/\alpha_b$. □

Next, we would like to prove an important invariant given that $\phi(G_b) \geq \alpha_b$: if $\mathsf{decomp}(H, I, \ell)$ is called, then $\mathsf{OPT}(I, \alpha_\ell) < \bar{s}_\ell$. In order to prove this, we need two lemmas.

Lemma 6.13. *Suppose that $\phi(G_b) \geq \alpha_b$. If $\mathsf{decomp}(H, I, 1)$ is called, then $\mathsf{OPT}(I, \alpha_1) < 4|A|/\alpha_b + 1 = \bar{s}_1$.*

Proof. Let $S'_1, \ldots, S'_{t'}$ be all the maximal sets as defined in the proof of Lemma 6.9. If $\mathsf{decomp}(H, I, 1)$ is called, we know that either $V_I \subset \bigcup_{i=1}^{t'} S'_i$ or $V_I \cap \bigcup_{i=1}^{t'} S'_i = \emptyset$.

First, if $V_I \subset \bigcup_{i=1}^{t'} S'_i$, then $|V_I| \leq |\bigcup_{i=1}^{t'} S'_i| \leq 4|A|/\alpha_b$ as shown in Claim 6.12. So $\mathsf{OPT}(I, \alpha_1) < 4|A|/\alpha_b + 1 = \bar{s}_1$. Second, if $V_I \cap \bigcup_{i=1}^{t'} S'_i = \emptyset$, let S be any α_1-sparse cut in I where $|S| \leq |V_I|/2$. We can similarly show, as in Lemma 6.9, that $\phi_G(V_I \cup \bigcup_{i=1}^{t'} S'_i) < \alpha_b/2$ and again have that $|V_I \cup \bigcup_{i=1}^{t'} S'_i| \leq 4|A|/\alpha_b$. So $\mathsf{OPT}(I, \alpha_1) < 4|A|/\alpha_b + 1 = \bar{s}_1$ again. □

Lemma 6.14. *For any subgraph $H = (V_H, E_H)$ of G induced by V_H and $\alpha' < \alpha_b/2$, let A_H is the set of endpoints of edges in $\partial_G(V_H)$ and $B_H = (A \cup A_H) \cap V_H$. If $\phi(G_b) \geq \alpha_b$, then any α'-sparse cut $S \subset V_H$ where $|S| \leq |V_H - S|$ must be (B_H, σ)-overlapping in H. That is, if $\phi(G_b) \geq \alpha_b$, then $\mathsf{OPT}(H, \alpha') = \mathsf{OPT}(H, \alpha', B_H, \sigma)$.*

Proof. First, consider any cut edge $(u,v) \in \partial_{G_b}(S)$ in the initial graph G_b where $u \in S$. We claim that either $(u,v) \in \partial_H(S)$ or $u \in A \cup A_H$. Indeed, if $u \notin A \cup A_H$, i.e. u is not incident to any edge in D nor $\partial_G(V_H)$, all edges incident to u is inside H, and hence $(u,v) \in \partial_H(S)$. Since there are at most $\Delta|S \cap (A \cup A_H)|$ edges incidents to all nodes $u \in S \cap (A \cup A_H)$. It follows that $\delta_{G_b}(S) \leq \delta_H(S) + \Delta|S \cap (A \cup A_H)|$.

Suppose that there is an α'-sparse cut $S \subset V_H$ is not (B_H, σ)-overlapping, i.e. $|S \cap (A \cup A_H)| < \sigma|S|$. Then we have that

$$\delta_H(S) \geq \delta_{G_b}(S) - \Delta|S \cap (A \cup A_H)| > \alpha_b|S| - \sigma\Delta|S| = \alpha_b|S| - \alpha_b|S|/2 = \alpha_b|S|/2.$$

That is, $\phi_H(S) \geq \alpha_b/2 > \alpha'$, which is a contradiction. □

Now, we can prove **the main invariant**:

Lemma 6.15. *Suppose that $\phi(G_b) \geq \alpha_b$. If $\mathsf{decomp}(H, I, \ell)$ is called, then the invariant $\mathsf{OPT}(I, \alpha_\ell) < \bar{s}_\ell$ is satisfied.*

Proof. When $\ell = 1$, $\mathsf{OPT}(I, \alpha_\ell) < \bar{s}_1$ by Lemma 6.13. In particular, the invariant is satisfied when $\mathsf{decomp}(G, G, 1)$ or $\mathsf{decomp}(H[S], H[S], 1)$ is called. The invariant $\mathsf{decomp}(H[V_H \setminus S], I, \ell)$ is the same as the one for $\mathsf{decomp}(H, I, \ell)$, and hence is satisfied by induction.

Finally, we claim that the invariant is satisfied when $\mathsf{decomp}(H, H, \ell{+}1)$ is called, i.e., $\mathsf{OPT}(H, \alpha_{\ell+1}) < \bar{s}_{\ell+1}$. By Step 5.a, $|S| < \bar{s}_{\ell+1}/c_{size}$. By Step 5, $\mathsf{OPT}(H, \alpha_{\ell+1}, B_H, \sigma)/c_{size} \leq |S|$ as $\alpha_{\ell+1} = \alpha_\ell/c_{exp}$. Since H is induced by V_H and $\alpha_{\ell+1} \leq \alpha_1 < \alpha_b/2$ satisfying the conditions in Lemma 6.14, we have $\mathsf{OPT}(H, \alpha_{\ell+1}) = \mathsf{OPT}(H, \alpha_{\ell+1}, B_H, \sigma)$. Therefore, $\mathsf{OPT}(H, \alpha_{\ell+1}) \leq c_{size}|S| < \bar{s}_{\ell+1}$ as desired. □

Finally, we bound the expansion of the components of G^d.

Lemma 6.16. *Suppose that $\phi(G_b) \geq \alpha_b$. Each connected component C in G^d either is a singleton or, otherwise, $\phi(C) \geq \alpha$.*

Proof. C can be returned as a connected component in G^d in either Step 2, 3 or 4 in Algorithm 6.1.

If C is returned in Step 2, then C is a connected component in a graph induced by E^d_H returned by $\mathcal{A}_{decomp}(H, \alpha, p')$. By Theorem 5.1, if C is not a singleton, then $\phi(C) \geq \alpha$.

If C is returned in Step 3, then we know $\mathsf{decomp}(H, H, B, L)$ was called, and $C = H$. By the invariant, we have $\mathsf{OPT}(H, \alpha_L) < \bar{s}_L \leq 1$, i.e. there is no α_L-sparse cut in H. As $\alpha_L = \alpha$, $\phi(H) \geq \alpha$.

If C is returned in Step 4, then we know $C = H$. If H is not a singleton, then $\mathcal{A}_{cut}(H, \alpha_\ell, B, \sigma)$ reports that there is no (α_ℓ/c_{exp})-sparse (B_H, σ)-overlapping cut in H where $\ell < L$. By Lemma 6.14, there is actually no (α_ℓ/c_{exp})-sparse cut in H, i.e. $\phi(H) \geq \alpha_\ell/c_{exp} \geq \alpha_L = \alpha$. □

We conclude the correctness of the algorithm from Lemma 6.16 and Lemma 6.9 the following:

Corollary 6.17. *Assume that $\mathcal{A}_{decomp}$ is deterministic. Suppose that $\phi(G_b) \geq \alpha_b$. $\mathsf{decomp}(G, G, 1)$ outputs E^s such that $|E^s| \leq 4\Delta|A|/\alpha_b$, and each connected component C of G^d is either a singleton or has high expansion: $\phi(C) \geq \alpha$.*

Now, it is left to analyze the running time.

6.2.3 Running time

In this section, we prove that if $\phi(G_b) \geq \alpha_b$, then $\mathsf{decomp}(G, G, 1)$ takes at most $\bar{t}$ time. In other words, if $\mathsf{decomp}(G, G, 1)$ takes more that $\bar{t}$ time, then $\phi(G_b) < \alpha_b$.

To analyze the running time, we need some more notation. We define the recursion tree $\mathcal{T}$ of $\mathsf{decomp}(G, G, 1)$ as follows:

- Each node of $\mathcal{T}$ represents the parameters of the procedure $\mathsf{decomp}(H, I, \ell)$.
- $(G, G, 1)$ is the root node.
- For each (H, I, ℓ), if $\mathsf{decomp}(H, I, \ell)$ returns a component either in Step 2,3, or 4 Algorithm 6.1, then (H, I, ℓ) is a leaf.
- If $\mathsf{decomp}(H, I, \ell)$ recurses on $\mathsf{decomp}(H[S], H[S], 1)$ and $\mathsf{decomp}(H[V_H \backslash S], I, \ell)$, then $(H[S], H[S], 1)$ and $(H[V_H \setminus S], I, \ell)$ is a left and right children of (H, I, ℓ) respectively, and the edges $((H, I, \ell), (H[S], H[S], 1))$ and $((H, I, \ell), (H[V_H \setminus S], I, \ell))$ are *left edge* and *right edge* respectively.
- If $\mathsf{decomp}(H, I, \ell)$ recurses on $\mathsf{decomp}(H, H, \ell + 1)$, then $(H, H, \ell + 1)$ is the only child of (H, I, ℓ), and the edge $((H, I, \ell), (H, H, \ell + 1))$ is a *down edge*.

Lemma 6.18. *Suppose that $\phi(G_b) \geq \alpha_b$. If the depth of the recursion tree $\mathcal{T}$ is D, then the total running time of $\mathsf{decomp}(G, G, 1)$ is $\tilde{O}(\frac{\Delta^{4.5}|A|^{1.5}}{\alpha_b^3}\gamma D \log \frac{1}{p})$ where $\gamma = n^{O(\sqrt{\log\log n/\log n})}$.*

Proof. The running time on $\mathsf{decomp}(H, I, \ell)$, excluding the time spent in the next recursion level, is contributed by either 1) the running time of $\mathcal{A}_{decomp}(H, \alpha, p')$ when $|V_H - B_H| < \frac{3}{\sigma}|B_H|$, 2) the running time of $\mathcal{A}_{cut}(H, \alpha_\ell, B_H, \sigma)$, and 3) the time for outputting H as the component of G^d when H is not the largest component.

The total running time of the third case is bounded by $O(\Delta|\bigcup_{i=1}^{t'} S_i'|) = O(\Delta|A|/\alpha_b)$ by Claim 6.12 where $\{S_i'\}_i$ are the *maximal* set defined in Lemma 6.9. It is left to bound the running time of the first and second cases.

In the first case, we have that $|E_H| \leq 2\Delta|V_H| = O(\frac{\Delta}{\sigma}|B_H|)$. So, by Theorem 5.1, the running time of $\mathcal{A}_{decomp}(H, \alpha, p')$ is $O(|E_H|\gamma \log \frac{1}{p'}) = \tilde{O}(\frac{\Delta}{\sigma}|B_H|\gamma \log \frac{1}{p})$ because $p' = p/n^2$ and $\gamma = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$. In the second case, the running time of $\mathcal{A}_{cut}(H, \alpha_\ell, B_H, \sigma)$ is $\tilde{O}((\Delta|B_H|/\sigma)^{1.5})$ by Theorem 6.5. In either case, the running time is at most $\tilde{O}((\Delta|B_H|/\sigma)^{1.5}\gamma \log \frac{1}{p})$.

Let $\mathcal{T}_d$ be the set of all the nodes in $\mathcal{T}$ of depth *exactly* d. Let A^s be the endpoints of edges in E^s. We can see that, for any two nodes $(H, I, \ell), (H', I', \ell') \in \mathcal{T}_d$, H and H' are disjoint subgraphs of G, and so B_H and $B_{H'}$ are disjoint subset of $A \cup A^s$. So, the total running time we spent on $\mathcal{T}_d$ is $\tilde{O}((\Delta|A \cup A^s|/\sigma)^{1.5}\gamma \log \frac{1}{p})$. Hence, the total running time of $\mathcal{T}$ is $\tilde{O}((\Delta|A \cup A^s|/\sigma)^{1.5}\gamma D \log \frac{1}{p})$. As $\sigma = \alpha_b/2\Delta$ and $|A^s| = O(\Delta|A|/\alpha_b)$ by Lemma 6.9, the running time is

$$\tilde{O}((\frac{\Delta(\Delta|A|/\alpha_b)}{\alpha_b/2\Delta})^{1.5}\gamma D \log \frac{1}{p}) = \tilde{O}(\frac{\Delta^{4.5}|A|^{1.5}}{\alpha_b^3}\gamma D \log \frac{1}{p}).$$

□

Now, we bound the depth D of $\mathcal{T}$. Recall $L \leq 1/\epsilon$, $c_{size} = 3/\sigma = 6\Delta/\alpha_b$ and $\bar{s}_1 = 4|A|/\alpha_b + 1$.

Lemma 6.19. *Suppose that $\phi(G_b) \geq \alpha_b$. Let P be any path from leaf to root of $\mathcal{T}$. P contains at most $\log n$ left edges, $L \log n$ down edges, and $L \log n \times k$ right edges where $k = c_{size}\bar{s}_1^\epsilon$. That is, the depth of $\mathcal{T}$ is $D = L \log n \times c_{size}\bar{s}_1^\epsilon = \tilde{O}(\frac{\Delta|A|^\epsilon}{\alpha_b^{1+\epsilon}})$.*

Proof. Consider the left edge $((H, I, \ell), (H[S], H[S], 1))$. Since $|S| \leq |V_H|/2$, P contains $\log n$ left edges. Between any two left edges in P there are at most L down edges because Step 3 in Algorithm 6.1 always terminates the recursion when $\ell = L$. So P contains at most $L \log n$ down edges. To prove that there are at most $L \log n \times k$ right edges in P, it suffices to prove that there cannot be γ right edges between any left edges or down edges in P.

Suppose that $(H_1, I, \ell), \ldots, (H_\gamma, I, \ell)$ are nodes in P where, for each i, $((H_i, I, \ell), (H_{i+1}, I, \ell))$ is a right edge, and (H_1, I, ℓ) is a deeper endpoint of left edges or a down edges (hence $I = H_1$).

For each i, let S_i be the cut such that $H_{i+1} = H_i[V_{H_i} \setminus S_i]$ and $\phi_{H_i}(S_i) < \alpha_\ell$. Since $\{S_i\}_{i \leq k}$ are mutually disjoint. We conclude $\phi_{H_1}(\bigcup_{i=1}^k S_i) < \alpha_\ell$ using the same argument as in Lemma 6.9. However, we also have that $|S_i| \geq \bar{s}_{\ell+1}/\bar{c}_{size}$, for all i, and hence $|\bigcup_{i=1}^k S_i| \geq k\bar{s}_{\ell+1}/\bar{c}_{size} \geq \bar{s}_1^\epsilon \bar{s}_{\ell+1} = \bar{s}_\ell$. So $\bigcup_{i=1}^k S_i$ contradicts the invariant for $\mathsf{decomp}(H_1, I, \ell)$, where $I = H_1$, which says $\mathsf{OPT}(H_1, \alpha_\ell) < \bar{s}_\ell$. Note that the invariant must hold by Lemma 6.15. □

Corollary 6.20. *Suppose that $\phi(G_b) \geq \alpha_b$. $\mathsf{decomp}(G, G, 1)$ runs in time $\bar{t} = \tilde{O}(\frac{\Delta^{5.5}|A|^{1.5+\epsilon}}{\alpha_b^{4+\epsilon}\epsilon}\gamma \log \frac{1}{p})$.*

Proof. We have $D = \tilde{O}(\frac{\Delta|A|^\epsilon}{\alpha_b^{1+\epsilon}\epsilon})$ by Lemma 6.19. By Lemma 6.18, we have that $\mathsf{decomp}(G, G, 1)$ runs in time

$$\tilde{O}(\frac{\Delta^{4.5}|A|^{1.5}}{\alpha_b^3}\gamma D \log \frac{1}{p}) = \tilde{O}(\frac{\Delta^{5.5}|A|^{1.5+\epsilon}}{\alpha_b^{4+\epsilon}\epsilon}\gamma \log \frac{1}{p})$$

. □

Now, we can conclude the main theorem.

Proof of Theorem 6.1. From Section 6.2.1, we have that in time $\bar{t} = \tilde{O}(\frac{\Delta^{5.5}|A|^{1.5}n^\epsilon}{\alpha_b^4\epsilon}\gamma)$ $\mathsf{decomp}(G, G, 1)$ outputs $E^s \subset E$ and all components, except the largest one, of G^d where $G^d = (V, E^d) = (V, E - E^s)$. Otherwise, by Corollary 6.20, we report $\phi(G_b) < \alpha_b$.

When we assume that $\mathcal{A}_{decomp}$ is deterministic. If $\phi(G_b) \geq \alpha_b$, by Corollary 6.17, we have $|E^s| \leq 4\Delta|A|/\alpha_b$ and each connected component C of G^d is either a singleton or has high expansion $\phi(C) \geq \alpha$. To remove the assumption that $\mathcal{A}_{decomp}$ is deterministic, and show that with probability $1 - p$, the outputted decomposition is correct. Observe that, as the depth of the recursion tree $\mathcal{T}$ of $\mathsf{decomp}(G, G, 1)$ is $D = L \log n \cdot c_{size}\bar{s}_1^\epsilon$, then $\mathcal{A}_{decomp}$ is called at most nD times. So the probability that $\mathcal{A}_{decomp}$ always works correctly is at least $1 - p'nD \geq 1 - p$. □

7 Monte Carlo Dynamic SF with $O(n^{0.4+o(1)})$ Update Time

In this section, the goal is to prove:

Theorem 7.1. *There is a randomized dynamic SF algorithm for any graphs with n nodes and m initial edges that works correctly against adaptive adversaries with high probability and has preprocessing time $O(m^{1+o(1)})$ and worst-case update time $O(n^{0.4+o(1)})$ where the term $o(1) = O(\sqrt{\log\log n / \log n})$.*

By Corollary 4.19, we only need to prove the lemma below. We note that the preprocessing time $t_p(n, p)$ below does *not* depend on the length of update sequence T.

Lemma 7.2. *There is a randomized dynamic SF algorithm for 3-bounded degree graphs with n nodes and $\Theta(n)$ edges that works correctly against adaptive adversaries with probability $1 - p$ for the first T updates and has preprocessing time $t_p(n, p) = O(n^{1+o(1)} \log \frac{1}{p})$ and worst-case update time*

$t_u(n,p,T)$ *such that* $t_u(n,p,n^{0.6}) = O(n^{0.4+o(1)}\log\frac{1}{p})$. *The term* $o(1) = O(\sqrt{\log\log n/\log n})$ *in both preprocessing and update time.*

Throughout this section, we let $\beta = \beta(n) = 2^{O(\log\log n)^3} = n^{o(1)}$ and $\gamma = \gamma(n) = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$ be the function depending on n from Theorem 4.13 and Theorem 5.1 respectively.

7.1 Dynamic SF on Expanders

The following lemma shows a decremental SF algorithm that runs on a graph which is initially an expander. It exploits a cut recovery tree from Section 4.2.

Lemma 7.3. *There is a decremental* SF *algorithm* $\mathcal{A}$ *for 3-bounded degree graphs with* n *nodes and has initial expansion at least* α*, that works correctly against adaptive adversaries with probability* $1-p$ *for the first* T *updates and* $\mathcal{A}$ *has preprocessing time* $\tilde{O}(n\beta)$ *and worst-case update time* $\tilde{O}((1/\alpha+\sqrt{T/\alpha})\beta\log\frac{1}{p})$.

Let $G=(V,E)$ be the initial graph. $\mathcal{A}$ preprocesses G as follows: 1) find any spanning tree F of G. Then initialize both an augmented ET tree $\mathcal{E}$ and a cut recovery tree $\mathcal{D}$ with parameter $k = \sqrt{T/\alpha}$ on (G,F). By Theorem 3.3 and Theorem 4.13, this takes time $O(|E|+|E|\beta\log n) = \tilde{O}(n\beta)$.

Next, $\mathcal{A}$ handles edge deletion as follows. If the deleted edge e is not in tree edge, then just remove e from G. Otherwise, if $e=(u,v)$ is a tree edge of F then we try to find a replacement edge e' for reconnecting the two separated components of F as described in Algorithm 7.1 and, if e' is returned, add e' into F. We set $\sigma = O((\frac{1}{\alpha}+\sqrt{\frac{T}{\alpha}})\log\frac{T}{p})$.

1. Using $\mathcal{E}$, set $T_u = \mathsf{find_tree}(u)$ and $T_v = \mathsf{find_tree}(v)$ where $\mathsf{tree_size}(T_u) < \mathsf{tree_size}(T_v)$.
2. Sample independently σ non-tree edges in $G\setminus F$ incident to T_u using $\mathcal{E}$. If there is a sampled edge e' from $\partial_G(V(T_u))$, then RETURN e'.
3. Query k-$\mathsf{cutset}(T_u)$ to $\mathcal{D}$. If the returned set $\tilde{\partial}\neq\emptyset$, then RETURN a random edge $e'\in\tilde{\partial}$.
4. RETURN "no replacement edge"

Algorithm 7.1: The algorithm for finding a replacement edge e' if exists, after deleting an edge $e=(u,v)\in F$.

Claim 7.4. Algorithm 7.1 runs in time $\tilde{O}((1/\alpha+\sqrt{T/\alpha})\beta\log\frac{1}{p})$.

Proof. Step 1 takes time at most $O(\log n)$. Step 2 takes time $O(\sigma\log n) = O((\frac{1}{\alpha}+\sqrt{\frac{T}{\alpha}})\log\frac{1}{p}\log n))$ by Theorem 3.3. Step 3 takes time $O(k\beta\log^2 n) = O(\sqrt{\frac{T}{\alpha}}\beta\log^2 n)$ by Theorem 4.13. So the running time in Algorithm 7.1 is at most $\tilde{O}((1/\alpha+\sqrt{T/\alpha})\beta\log\frac{1}{p})$. □

Claim 7.5. Algorithm 7.1 returns a replacement edge, if exists, with probability $1-p$.

Proof. Let $S = V(T_u)$ where T_u is defined in Algorithm 7.1. If a replacement edge exists, then $\delta_G(S)\geq 1$. If $1\leq\delta_G(S)\leq\sqrt{T/\alpha}$, then by the cut recovery tree $\mathcal{D}$, all cut edges $\partial_G(S)$ is returned, and hence so is a replacement edge.

Otherwise, if $\delta_G(S) > \sqrt{T/\alpha}$, then we claim that $\phi_G(S)\geq\min\{\frac{\alpha}{2},\frac{1}{2\sqrt{T/\alpha}}\}$. There are two cases. For the first case, if $|S|\geq 2T/\alpha$, then we claim that $\phi_G(S)\geq\alpha/2$. Indeed, let G_0 be an initial graph.

We have that $\delta_G(S) \geq \delta_{G_0}(S) - T \geq \alpha|S| - T \geq \alpha|S| - \frac{\alpha}{2}|S| = \frac{\alpha}{2}|S|$ where the third inequality is because $\phi(G_0) \geq \alpha$. For the second case, we have $|S| < 2T/\alpha$ while $\delta_G(S) > \sqrt{T/\alpha}$ and therefore $\phi_G(S) \geq \frac{\sqrt{T/\alpha}}{2T/\alpha} = \frac{1}{2\sqrt{T/\alpha}}$. By the property of the augmented ET-tree $\mathcal{E}$, the probability that a sampled edge is from $\partial_G(S)$ is $\frac{\delta_G(S)}{nt_vol_{(G,F)}(S)} \geq \frac{\delta_G(S)}{3|S|} \geq \frac{\phi_G(S)}{3}$. Now, the probability that all sampled σ edges are not from $\partial_G(S)$ is $(1 - \frac{\phi_G(S)}{3})^\sigma \leq (1 - \frac{\max\{\frac{\alpha}{2}, \frac{1}{2\sqrt{T/\alpha}}\}}{3})^\sigma \leq p$, by choosing the constant in $\sigma = O((\frac{1}{\alpha} + \sqrt{\frac{T}{\alpha}}) \log \frac{1}{p})$. □

Now, we conclude the proof of Lemma 7.3.

Proof of Lemma 7.3. The preprocessing time of $\mathcal{A}$ is $\tilde{O}(n\beta)$. It is clear that for each update, the bottleneck is the running time of Algorithm 7.1 for finding replacement edge which takes $\tilde{O}((1/\alpha + \sqrt{T/\alpha})\beta \log \frac{1}{p})$ by Claim 7.4, and each update $\mathcal{A}$ correctly find a replacement edge with probability $1 - p$ by Claim 7.5. □

7.2 Reduction to Expanders via Global Expansion Decomposition

The following reduction shows that, given a dynamic SF algorithm that runs on expanders, we can obtain another algorithm that needs not run on expanders. The main idea is to use the global expansion decomposition algorithm from Theorem 5.1 to partition edges on the graph into a collection of expanders and some small set of the remaining edges. We then separately maintain a spanning tree in each expander using the algorithm from Lemma 7.3, and "combine" them with the remaining edges using 2-dim ET tree from Theorem 4.2.

For any n, α, p, and T, for technical reason, let $t_p(n, \alpha, p, T)$ and $t_u(n, \alpha, p, T)$ be some functions such that, for any n_1, n_2 where $n_1 + n_2 \leq n$, we have $t_p(n_1, \alpha, p, T) + t_p(n_2, \alpha, p, T) \leq t_p(n, \alpha, p, T)$ and $t_u(n_1, \alpha, p, T) \leq t_u(n, \alpha, p, T)$.

Lemma 7.6. *For any n, α, p, and T, suppose that there is a decremental SF algorithm $\mathcal{A}$ for 3-bounded degree graphs with n nodes and initial expansion at least α, that works correctly against adaptive adversaries with probability $1 - p$ for the first T updates and $\mathcal{A}$ has preprocessing time $t_p(n, \alpha, p, T)$ and worst-case update time $t_u(n, \alpha, p, T)$. Then there is a dynamic SF algorithm $\mathcal{B}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with probability $1 - O(nTp)$ for the first T updates and $\mathcal{B}$ has preprocessing time $O(t_p(n, \alpha, p, T) + n\gamma \log \frac{1}{p})$ and worst-case update time $O(t_u(n, \alpha, p, T) + \sqrt{n\gamma\alpha + T})$.*

Before proving Lemma 7.6, we show that it implies Lemma 7.2.

Proof of Lemma 7.2. We set $\alpha = 1/n^{0.2}$. By Lemma 7.3 and Lemma 7.6, there is dynamic SF algorithm $\mathcal{B}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with probability $1 - p$ for the first T updates and $\mathcal{B}$ has preprocessing time $t_p(n, p) = \tilde{O}(n\beta + n\gamma \log \frac{1}{p}) = O(n^{1+o(1)} \log \frac{1}{p})$ and worst-case update time $t_u(n, p, T) = \tilde{O}((1/\alpha + \sqrt{T/\alpha})\beta \log \frac{1}{p} + \sqrt{n\gamma\alpha + T})$. We can verify that $t_u(n, p, n^{0.6}) = O(n^{0.4+o(1)} \log \frac{1}{p})$ and the term $o(1) = O(\sqrt{\log \log n / \log n})$ in both preprocessing and update time. □

Now we prove Lemma 7.6. Given the algorithm $\mathcal{A}$ that runs on a graph with initial expansion at least α by the assumption, for convenience, let $\mathcal{A}'$ be the algorithms that runs on a graph whose all connected components has initial expansion at least α. $\mathcal{A}'$ works by simply running $\mathcal{A}$ on each component.

The target algorithm $\mathcal{B}$ preprocesses a graph $G = (V, E)$ as described in Algorithm 7.2. The idea is to partition edges E into E^d and E^s using the global expansion decomposition algorithm from Theorem 5.1 with success probability $1 - p$. By adding the failure probability of $\mathcal{B}$ by p at the end of analysis, we now assume that the decomposition succeeds. Then let F^d be a spanning forest of $G^d = (V, E^d)$ maintained by $\mathcal{A}'$. The key invariant to maintain a spanning forest F of G such that $F^d \subseteq F$.

1. Run the expansion decomposition on G with expansion parameter α and failure probability p and obtain $G^d = (V, E^d)$ and $G^s = (V, E^s)$.
2. Preprocess G^d using $\mathcal{A}'$ and obtain a spanning forest F^d.
3. Construct a spanning forest F of G where $F \supseteq F^d$.
4. Preprocess $(F \cup E^s, F)$ using a 2-dim ET tree $\mathcal{D}$.

Algorithm 7.2: The algorithm for preprocessing the initial graph $G = (V, E)$

Claim 7.7. The preprocessing time of $\mathcal{B}$ is $O(t_p(n, \alpha, p, T) + n\gamma \log \frac{1}{p})$.

Proof. By Theorem 5.1, Step 1 takes $O(n\gamma \log \frac{1}{p})$. Let $n_1, \ldots, n_k$ be the size of all connected components in G^d. We have $\sum_{i=1}^{k} n_i \leq n$, so Step 2 takes $\sum_{i=1}^{k} t_p(n_i, \alpha, p, T) \leq t_p(n, \alpha, p, T)$. Step 3 can be done easily in $O(n)$. Step 4 takes in $O(n + |E^s|) = O(n)$ time by Theorem 4.2. □

Remark 7.1. In the following description of how to handle the updates, we maintain the following invariants: 1) the underlying graph and forest of $\mathcal{D}$ is $(F \cup E^s, F)$ and 2) $\mathcal{A}'$ maintains a dynamic forest F^d on the underlying graph G^d. Therefore, so whenever G^d, E^s, F^d or F are changed, we feed the update to $\mathcal{A}'$ and $\mathcal{D}$ accordingly.

To insert $e = (u, v)$, we always insert e into E^s, and if u and v are not connected in F (this can be checked using the operation find_tree of the 2-dim ET tree $\mathcal{D}$), then we also add e into F.

To delete $e = (u, v)$, if e was a tree-edge, then we find a replacement edge e' for reconnecting the two separated components of F as described in Algorithm 7.3 and, if e' is returned, add e' into F.

1. If $e \in F^d$ then feed the deletion to $\mathcal{A}'$. If $\mathcal{A}'$ returns a replacement edge $e' \in E^d$ of F^d, then RETURN e'.
2. Using $\mathcal{D}$, set $T_u = \mathsf{find_tree}(u)$ and $e' = \mathsf{find_edge}(T_u)$.
3. If $e' \neq \emptyset$, then RETURN e'. Else RETURN "no replacement edge"

Algorithm 7.3: The algorithm for finding a minimum weight replacement edge e' if exists, after deleting an edge $e = (u, v) \in F$.

Claim 7.8. F is a spanning forest of G such that $F \supseteq F^d$ throughout the whole update sequences with probability $1 - pnT$.

Proof. We first assume that $\mathcal{A}'$ is deterministic and we will remove it later. We prove that F is a spanning forest of G where $F \supseteq F^d$ after each edge deletion.

It is clear that $F \supseteq F^d$ because, after preprocessing, we have $F \supseteq F^d$ and then, in Step 1 of Algorithm 7.3, whenever $\mathcal{A}'$ adds an edge e' into F^d, we always add e' into F. We only need to argue that F is a spanning forest of G. Indeed, F is a forest in G because an edge (u, v) is added into F only if u and v are not connected in F.

Now, we claim that F spans G. It suffices to prove that after each deletion if a replacement edge exists, then Algorithm 7.3 can find one. Suppose that $e = (u, v) \in F$ is deleted. Consider the first case where $e \notin F^d$. Let $T_u \cup (u, v) \cup T_v$ be a connected component in F before deletion. Any replacement edge e', if exists, must be in $E^s(T_u, T_v)$, i.e. a set of edges $(x, y) \in E^s$ where $x \in V(T_u)$ and $y \in V(T_u)$). Indeed, $e' \notin E^d$ because, after deletion, F^d still spans all components in $G^d = (V, E^d)$. So $e' \in E^s$, and moreover $e' \in E^s(T_u, T_v)$ otherwise e' cannot reconnect T_u and T_v. Now, Step 2 of Algorithm 7.3 will return an edge from $E^s(T_u, T_v)$. Consider the second case where $e \in F^d$. Let $T_u^d \cup (u, v) \cup T_v^d$ be a connected component in F^d. Any replacement edge e', if exists, must be in $E^d(T_u^d, T_v^d) \cup E^s(T_u, T_v)$ for the similar reason. An edge $E^d(T_u^d, T_v^d)$ and $E^s(T_u, T_v)$ will be found in Step 1 and Step 2 of Algorithm 7.3 respectively.

Finally, we remove the assumption that $\mathcal{A}$ is deterministic. As $\mathcal{A}$ works correctly at each step with probability $1 - p$, and $\mathcal{A}'$ runs instances of $\mathcal{A}$ on at most n components of G^d for T time steps, so $\mathcal{A}'$ works correctly throughout the whole update sequences with probability at least $1 - pnT$. □

Claim 7.9. The update time of $\mathcal{B}$ is $O(t_u(n, \alpha, p, T) + \sqrt{n\gamma\alpha + T})$.

Proof. Observe the the bottleneck of the update time is the update time of $\mathcal{A}'$ and 2-dim ET tree $\mathcal{D}$. Let $n_1, \ldots, n_k$ be the size of all components in G^d. The update time of $\mathcal{A}'$ is at most $\max_{i=1}^{k} t_u(n_i, \alpha, p, T) \leq t_u(n, \alpha, p, T)$. Let E_0^s be the set E^s right after preprocessing. We have that $E^s \subset E_0^s \cup I$ where I is the set of inserted edges and we know $|I| \leq T$ as there are at most T updates. As E^s is the set of non-tree edges in the underlying graph of $\mathcal{D}$, each operation of $\mathcal{D}$ can be done in time $O(\sqrt{|E^s|}) = O(\sqrt{|E_0^s \cup I|}) = O(\sqrt{n\gamma\alpha + T})$ by Theorem 5.1. So the total update time is $O(t_u(n, \alpha, p, T) + \sqrt{n\gamma\alpha + T})$. □

We conclude the proof of Lemma 7.6.

Proof of Lemma 7.6. Given $\mathcal{A}$, we have that $\mathcal{B}$ has preprocessing time $O(t_p(n, \alpha, p, T) + n\gamma \log \frac{1}{p})$ by Claim 7.7, and update time $O(t_u(n, \alpha, p, T) + \sqrt{n\gamma\alpha + T})$ by Claim 7.9. By Claim 7.8, $\mathcal{B}$ works correctly for *all* updates with probability $1 - pnT$. In particular, $\mathcal{B}$ works correctly for each update with the same probability. □

7.2.1 Reduction for Las Vegas algorithms

Since the reduction in Lemma 7.6 is only for Monte Carlo algorithm, next, we state a similar reduction for Las Vegas algorithm which we will use in Section 8.

Lemma 7.10. *For any n, α, p and T, suppose that there is a decremental* SF *algorithm $\mathcal{A}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with certainty for the first T updates. Moreover, suppose that $\mathcal{A}$ has the following properties: Given an initial graph G_b, $\mathcal{A}$ takes $t_p(n, \alpha, p, T)$ time to preprocess or report failure. Given each edge deletion, $\mathcal{A}$ takes $t_p(n, \alpha, p, T)$ worst-case time to update or report failure. If $\phi(G_b) \geq \alpha$, then, after preprocessing and for each update, $\mathcal{A}$ does not report failure with probability $1 - p$.*

Then there is a dynamic SF *algorithm $\mathcal{B}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with certainty for the first T updates and $\mathcal{B}$ has preprocessing time $O(t_p(n, \alpha, p, T) + n\gamma \log \frac{1}{p})$ and worst-case update time $O(t_u(n, \alpha, p, T) + \sqrt{n\gamma\alpha + T})$ with probability $1 - O(p)$.*

Lemma 7.10 is different from Lemma 7.6 as follows: 1) $\mathcal{A}$ and $\mathcal{B}$ work correctly with certainty, 2) the initial graph G_b of $\mathcal{A}$ may not have expansion α, but only when $\phi(G_b) \geq \alpha$, then $\mathcal{A}$ will not fail with high probability. Since the reduction works in very similar way as in Lemma 7.6, we will just sketch the proof below.

Proof sketch. We use the same algorithm as in Lemma 7.6. Given any graph $G = (V, E)$, the algorithm $\mathcal{B}$ first partitions edges E into E^d and E^s using the global expansion decomposition algorithm from Theorem 5.1 with success probability $1 - p$. We use the algorithm $\mathcal{A}$ to maintain a spanning tree on each component of $G^d = (V, E^d)$ as in Lemma 7.6. Whenever $\mathcal{A}$ reports failure, we just restart Algorithm 7.2 and obtain a new expansion decomposition. Since the decomposition succeeds with probability $1 - p$, the initial underlying graph of $\mathcal{A}$, which is a component C of G^d, has expansion at least α with probability $1 - p$. If $\phi(C) \geq \alpha$, $\mathcal{A}$ does not fails with probability $1 - p$ for each update. Therefore, with probability $1 - O(p)$, $\mathcal{B}$ has preprocessing time and update time as claimed, using the same reduction as in Lemma 7.6. □

8 Las Vegas Dynamic SF with $O(n^{0.49305+o(1)})$ Update Time

Throughout this section, let $\epsilon = 0.0069459$ and $\epsilon' = \sqrt{\frac{2\epsilon}{3\epsilon+0.5}} = 0.16332$. The main result of this section is the following:

Theorem 8.1. *There is a randomized dynamic* SF *algorithm for any graphs with n nodes and m initial edges that works correctly against adaptive adversaries with certainty and has preprocessing time $O(m^{1+o(1)})$ and worst-case update time $O(n^{1/2-\epsilon+o(1)})$ with high probability. The term $o(1) = O(\sqrt{\log\log n / \log n})$.*

To this end, it is enough the prove the following. Recall the factor $\gamma = n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$ from Theorem 5.1

Lemma 8.2. *For any n, p and T, there is a decremental* SF *algorithm $\mathcal{A}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with certainty for the first T updates. $\mathcal{A}$ preprocesses an initial graph G_b in time $t_p(n,p) = O(n)$. Then, given each edge deletion, $\mathcal{A}$ takes $t_u(n,p,T)$ worst-case time to update or report failure, where $t_u(n,p,n^{1/2+\epsilon}) = O(n^{1/2-\epsilon}\gamma \log \frac{1}{p})$. If $\phi(G_b) \geq 1/n^{2\epsilon}$, then, for each update, $\mathcal{A}$ does not report failure with probability $1 - O(p)$.*

Proof of Theorem 8.1. We set the expansion parameter in Lemma 7.10 as $\alpha_b = 1/n^{2\epsilon}$. Plugging Lemma 8.2 into Lemma 7.10, there is a dynamic SF algorithm $\mathcal{B}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with certainty for the first T updates and $\mathcal{B}$ has preprocessing time $t'_p(n,p) = O(n + n\gamma \log \frac{1}{p}) = O(n\gamma \log \frac{1}{p})$ and worst-case update time $t'_u(n,p,T) = O(t_u(n,p,T) + \sqrt{n\gamma\alpha_b + T})$ where $t'_u(n,p,n^{1/2+\epsilon}) = O(n^{1/2-\epsilon}\gamma \log \frac{1}{p})$ with probability $1 - O(p)$. Then, by plugging $\mathcal{B}$ into Corollary 4.20, we obtain Theorem 8.1. □

The rest of this section is to prove Lemma 8.2. Throughout this section, for any n, we also denote $U_b = n^{1/2+\epsilon}$ and $\alpha_b = 1/n^{2\epsilon}$. Let U_a be such that $U_a/\alpha_b^{1/\epsilon'} = n^{1/2-\epsilon}$ so $U_a = n^{1/2-\epsilon}\alpha_b^{1/\epsilon'} = n^{1/2-\epsilon-2\epsilon/\epsilon'}$.

8.1 Dynamic SF on Broken Expanders via Local Expansion Decomposition

The goal is the main technical lemma below, and then we will show later in Section 8.2 that the lemma below implies Lemma 8.2 by using a standard technique of maintaining two concurrent data

structures as in Remark B.1. Recall the definition of augmented ET tree and 2-dim ET tree from Definition 3.2 and Definition 4.1 respectively.

Lemma 8.3. *There is an algorithm $\mathcal{A}$ that can do the following:*

- *$\mathcal{A}$ is given $n, p, G_b, \alpha_b, D, F_0$, $\mathcal{E}$ and $\mathcal{E}^{2d}$ as inputs: $G_b = (V, E_b)$ is a 3-bounded degree graphs with n nodes. p is a failure probability parameter. α_b is an expansion parameter. $D \subset E_b$ is set of edges in G_b of size at most U_b. Let $G = (V, E_b - D)$ that $\mathcal{A}$ will compute the decomposition on. F_0 is some spanning forest of G. $\mathcal{E}$ is an augmented ET tree initialized on (G, F_0). $\mathcal{E}^{2d}$ is an 2-dim ET tree initialized on (F_0, F_0).*
- ***Preprocessing:** After given the input, $\mathcal{A}$ takes $O((\frac{U_b^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}})\log\frac{1}{p})$ time to either*
  1. *reports failure, or*
  2. *outputs the list of edges to be added or removed from F_0 to obtain another spanning forest F of G.*
- *Then, there is an online sequence of edge deletions of length U_a given to $\mathcal{A}$.*
- ***Update:** After each edge deletion, $\mathcal{A}$ takes $O(\frac{U_a}{\alpha_b^{1/\epsilon'}}\log\frac{1}{p} + \sqrt{\frac{U_b}{\alpha_b}}) = O(n^{1/2-\epsilon}\log\frac{1}{p})$ worst-case time and either*
  1. *reports failure, or*
  2. *outputs the list of edges to be added or removed from F to obtain new one spanning forest of G.*
- *If $\phi(G_b) \geq \alpha_b$, then, after preprocessing and for each update, $\mathcal{A}$ does not report failure with probability $1 - p$.*

We call the time $\mathcal{A}$ takes before the update sequences *preprocessing time*, and the time $\mathcal{A}$ needs for each edge deletion *update time*.

Remark 8.1. In Lemma 8.3, if $D = \emptyset$, then we can substitute U_b by 0. This implies that $\mathcal{A}$ takes no time to preprocess (and without failure), and for each update in the sequence of length U_a, $\mathcal{A}$ takes $O(\frac{U_a}{\alpha_b^{1/\epsilon'}}\log\frac{1}{p}) = O(n^{1/2-\epsilon}\log\frac{1}{p})$ time to either report failure or update the spanning forest.

The intuition of Lemma 8.3 is clearer when we assume that the "before" graph G_b has expansion $\phi(G_b) \geq \alpha_b$. Then the initial graph G_0 is a "broken" expander, i.e. $G_0 = G_b - D$ is an expander with some edge deleted. Lemma 8.3 says that, with high probability, the algorithm $\mathcal{A}$ can preprocess in *sublinear time*, and obtain a dynamic SF algorithm that can handle U_a more deletions with update time $O(n^{1/2-\epsilon})$. The most important point of the algorithm $\mathcal{A}$ is that the preprocessing time is sublinear. This can be done by exploiting the local expansion decomposition algorithm from Section 6.

Now, we are ready to describe the algorithm for Lemma 8.3.

8.1.1 Preprocessing

Recall that the inputs are $n, p, G_b, \alpha_b, D, F_0$, $\mathcal{E}$ and $\mathcal{E}^{2d}$. Let $G = (V, E) = (V, E_b - D)$. We denote $(G_{\mathcal{E}}, F_{\mathcal{E}})$ and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}})$ the underlying graphs and forests of $\mathcal{E}$ and $\mathcal{E}^{2d}$ respectively. In the beginning, we are given $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G, F_0)$ and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (F_0, F_0)$.

First, we run the local expansion decomposition algorithm $\mathcal{A}_{local}$ from Theorem 6.1 with input $(n, p, 3, G_b, D, \alpha_b)$ and parameter ϵ'. If $\mathcal{A}_{local}$ reports that G_b has expansion less than α_b, then we also report failure and halt. Otherwise, $\mathcal{A}_{local}$ return the set of edges $E^s \subset E$ in G and the all components of G^d, except the largest one, where $G^d = (V, E^d) = (V, E - E^s)$. By Theorem 6.1, if $\phi(G_b) \geq \alpha_b$, then with probability $1 - p$, each connected components in G^d has expansion at least $\alpha = (\alpha_b/18)^{1/\epsilon'}$. Let $\sigma = O(\frac{1}{\alpha} \log \frac{1}{p})$. Throughout the sequence of edge deletions, we will maintain the following invariant: $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.

The preprocessing algorithm as described in Algorithm 8.1 is simply to validate these invariants before the first update comes.

1. If $\mathsf{count_tree}() > U_b$ on $F_{\mathcal{E}}$, then REPORT FAIL.
2. Run the local expansion decomposition algorithm $\mathcal{A}_{local}$ with input $(n, p, 3, G_b, D, \alpha_b)$ and parameter ϵ'. If $\mathcal{A}_{local}$ reports that $\phi(G_b) < \alpha_b$, then REPORT FAIL. Else, $\mathcal{A}_{local}$ returns $E^s \subset E$ in G_0 and all components, except the largest one, of $G^d = (V, E - E^s)$.
3. For each $e \in E^s$: //To update $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F_0 - E^s)$ and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F_0, F_0 - E^s)$
   (a) If $e \in F_0$, then $\mathsf{delete_tree}(e)$ on $F_{\mathcal{E}}$ and $F_{\mathcal{E}^{2d}}$.
   (b) $\mathsf{delete_nontree}(e)$ on $G_{\mathcal{E}}$ and $\mathsf{insert_nontree}(e)$ on $G_{\mathcal{E}^{2d}}$
4. While there is a component $T \in F_{\mathcal{E}}$ which not marked *big*: //To update $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F^d)$ and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F^d)$ where F^d is a spanning forest of G^d
   (a) Sample independently σ edges incident to T in $G_{\mathcal{E}}$.
   (b) If there is a sampled edge e' from $\partial_{G_{\mathcal{E}}}(V(T))$, then $\mathsf{insert_tree}(e')$ into $F_{\mathcal{E}}$ and $F_{\mathcal{E}^{2d}}$. Otherwise, REPORT FAIL.
   (c) Let $T' \in F_{\mathcal{E}}$ be the new component containing e'. Let C' be the component of G^d containing T'.
   (d) If $|V(T')| > |V(C')|/2$, then mark T' as *big*.
5. While there is a component $T \in F_{\mathcal{E}^{2d}}$ where $\mathsf{find_edge}(T) \neq \emptyset$: //To update $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G
   (a) Set $e' = \mathsf{find_edge}(T)$.
   (b) $\mathsf{insert_tree}(e')$ into $F_{\mathcal{E}^{2d}}$.

Algorithm 8.1: The algorithm for preprocessing the initial graph $G = (V, E)$

Now, we analyze Algorithm 8.1.

Claim 8.4. The running time of Algorithm 8.1 is $\tilde{O}((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}}) \log \frac{1}{p})$.

Proof. We will prove that even if Algorithm 8.1 does not fail, the running time is at most $\tilde{O}((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}}) \log \frac{1}{p})$.

Step 2 takes $\tilde{O}(\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma \log \frac{1}{p})$ by Theorem 6.1. If Algorithm 8.1 proceeds to Step 3, we have $|E^s| = O(U_b/\alpha_b)$ by Theorem 6.1. Step 3 takes $|E^s| \times O(\log n + \sqrt{|E^s|}) = O((U_b/\alpha_b)^{1.5})$ by Theorem 3.3 and Theorem 4.2 because as there are at most $|E^s|$ non-tree edges in $G_{\mathcal{E}^{2d}}$,

In Step 4, observe that in every iteration of the while loop the number of components T in $F_{\mathcal{E}}$ which is not marked as big always decreases by one. Right before Step 4, there are at most $|E^s|+U_b$ such components because $F_{\mathcal{E}}$ has at most U_b components after Step 1 and then there are at most $|E^s|$ tree edge deletions in Step 3. So there are at most $|E^s| + U_b$ iterations. In each iteration of the while loop in Step 4, it takes at most $\tilde{O}(\sigma) = \tilde{O}(\frac{1}{\alpha} \log \frac{1}{p}) = \tilde{O}(\frac{1}{\alpha_b^{1/\epsilon'}} \log \frac{1}{p})$ in Step 4.a. Step 4.b takes at most $O(\sqrt{|E^s|}) = O(\sqrt{U_b/\alpha_b})$ by Theorem 4.2. Note that Step 4.c and 4.d can be done in constant time because of the following. Using $\mathcal{E}$, we get $|V(T')| = \mathsf{tree_size}(T)$. We can also compute $|V(C')|/2$ where C' be the component of G^d containing T'. Let $u \in V(T)$. By Remark 6.1, we know a component C' of G^d where $u \in V(C)$, and also the size $|V(C')|$. In total, Step 4 takes $(|E^s| + U_b) \times \tilde{O}(\frac{1}{\alpha_b^{1/\epsilon'}} \log \frac{1}{p} + \sqrt{\frac{U_b}{\alpha_b}}) = O(\frac{U_b}{\alpha_b^{1+1/\epsilon'}} \log \frac{1}{p} + \frac{U_b^{1.5}}{\alpha_b^{1.5}})$.

In Step 5, again, in every iteration, the number of components T in $F_{\mathcal{E}^{2d}}$ decreases by one. Note that after Step 1, $F_{\mathcal{E}^{2d}} = F_{\mathcal{E}} = F_0$. By the same argument as we used for analyzing Step 4, there are at most $|E^s| + U_b$ iterations. Each iteration takes at most $O(\sqrt{|E^s|}) = O(\sqrt{U_b/\alpha_b})$. The total running time is $(|E^s| + U_b) \times O(\sqrt{\frac{U_b}{\alpha_b}}) = O(\frac{U_b^{1.5}}{\alpha_b^{1.5}})$.

We conclude that the total running time of Algorithm 8.1 is $\tilde{O}((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}}) \log \frac{1}{p})$. □

Next, we show that Algorithm 8.1 either reports failure or validate the invariants.

Claim 8.5. If Algorithm 8.1 does not report failure, then $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.

Proof. Suppose there is no failure. It is clear that $(G_{\mathcal{E}}, F_{\mathcal{E}})$ is updated to become $(G^d, F_0 - E^s)$ and (G^d, F^d) in Step 3 and 4 respectively. $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}})$ is updated to become $(E^s \cup F_0, F_0 - E^s)$, $(E^s \cup F^d, F^d)$, and $(E^s \cup F^d, F)$ in Step 3, 4, and 5 respectively. □

Claim 8.6. If $\phi(G_b) \geq \alpha_b$, then Algorithm 8.1 report failure with probability at most $O(p)$.

Proof. In Step 1, if $\mathsf{count_tree}() > U_b$ on $F_{\mathcal{E}}$, then in means that $G = (V, E_b - D)$ has more than U_b connected components and so G_b is not connected, i.e. $\phi(G_b) = 0$ which contradicts the assumption that $\phi(G_b) \geq \alpha_b$. In Step 2, the algorithm cannot fail because $\mathcal{A}_{local}$ reports failure only when $\phi(G_b) < \alpha_b$ by Theorem 6.1.

In Step 4.a, we sample σ edges incident to $T \in F_{\mathcal{E}}$ which is not marked as big, i.e. $|V(T)| \leq |V(C)|/2$ where C is the component of G^d containing T. As $\phi(G_b) \geq \alpha_b$, we have $\phi(C) \geq \alpha$ with probability $1 - O(p)$ by Theorem 6.1. Assuming that $\phi(C) \geq \alpha$, we have $\phi_C(V(T)) \geq \alpha$ and the probability that none of σ sampled edges is a cut edge is at most $(1 - \frac{\delta_C(V(T))}{vol_C(V(T))})^\sigma \leq (1 - \frac{\alpha}{3})^\sigma \leq p$. By union bound, the algorithm fails with probability at most $O(p)$. □

From the claims above can be summarized as follows:

Lemma 8.7. *After given the inputs, $\mathcal{A}$ takes $O((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}}) \log \frac{1}{p})$ time to either*

1. *reports failure, or*

2. *outputs the list of edges to be added or removed from F_0 to obtain another spanning forest F of G. Moreover, the underlying graphs and forested of the given augmented ET tree $\mathcal{E}$ and 2-dim ET tree $\mathcal{E}^{2d}$ are updated such that $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.*

Moreover, if $\phi(G_b) \geq \alpha_b$, $\mathcal{A}$ does not report failure with probability $1 - O(p)$.

Lastly, let us state the lemma in the corner case when $D = \emptyset$. In this case, we claim that the invariant already holds even if we do nothing. By setting $E^s = \emptyset$ and $F^d = F_0$, we have

Lemma 8.8. *After given the inputs where $D = \emptyset$, the underlying graphs and forested of the given augmented ET tree $\mathcal{E}$ and 2-dim ET tree $\mathcal{E}^{2d}$ are already such that $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G_b, F_0) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (F_0, F_0) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.*

8.1.2 Handling Updates

Throughout the update sequence, we keep the invariant, which is valid after preprocessing, that $(G_{\mathcal{E}}, F_{\mathcal{E}}) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.

The algorithm handles each edge deletion as standard: if the deleted edge $e \notin F$, then just remove e from G. That is, if $e \in E^d$, $\mathsf{delete_nontree}(e)$ from $G_{\mathcal{E}}$, otherwise $e \in E^s$, then $\mathsf{delete_nontree}(e)$ from $G_{\mathcal{E}^{2d}}$. Now, if $e \in F$, then we try to find a replacement edge e' for reconnecting the two separated components of F as described in Algorithm 8.2. Suppose that e' is returned. If $e' \in E^s$, then $\mathsf{insert_tree}(e')$ to $F_{\mathcal{E}^{2d}}$. If $e' \in E^d$, then $\mathsf{insert_tree}(e')$ to $F_{\mathcal{E}}$, $\mathsf{insert_nontree}(e')$ to $G_{\mathcal{E}^{2d}}$, and $\mathsf{insert_tree}(e')$ to $F_{\mathcal{E}^{2d}}$. Because whenever we add an edge to $F_{\mathcal{E}^{2d}}$ then we also add to $F_{\mathcal{E}}$, we conclude:

Fact 8.9. *The invariant $F \supseteq F^d$ is maintained.*

The important part is how Algorithm 8.2 works. Recall $\sigma = O(\frac{1}{\alpha} \log \frac{1}{p})$.

1. If $e \in F^d$ then
   (a) Using $\mathcal{E}$, set $T_u^d = \mathsf{find_tree}(u)$ and $T_v^d = \mathsf{find_tree}(v)$.
   (b) If $\mathsf{tree_size}(T_u^d) < \mathsf{tree_size}(T_v^d)$, then set $T^d = T_u^d$, otherwise $T^d = T_v^d$.
   (c) If $\mathsf{tree_size}(T^d) > 2U_a/\alpha$, then sample independently σ edges in G_d incident to T^d. If there is a sampled edge e' from $\partial_{G^d}(V(T^d))$, then RETURN e'. Otherwise, REPORT FAIL.
   (d) Else $\mathsf{tree_size}(T^d) \leq 2U_a/\alpha$, then list all edges in G_d incident to T^d. If there is a listed edge e' from $\partial_{G^d}(V(T^d))$, then RETURN e'.
2. Using $\mathcal{E}^{2d}$, set $T_u = \mathsf{find_tree}(u)$ and $e' = \mathsf{find_edge}(T_u)$.
3. If $e' \neq \emptyset$, then RETURN e'. Else RETURN "no replacement edge"

Algorithm 8.2: The algorithm for finding a minimum weight replacement edge e' if exists, after deleting an edge $e = (u, v) \in F$.

Claim 8.10. Algorithm 8.2 runs in time $O(\frac{U_a}{\alpha}\log\frac{1}{p}+\sqrt{\frac{U_b}{\alpha_b}}) = O(n^{1/2-\epsilon}\log\frac{1}{p}+n^{1/4+1.5\epsilon}) = O(n^{1/2-\epsilon}\log\frac{1}{p})$.

Proof. We prove that even there is no failure reported. Algorithm 8.2 runs in time $O(\frac{U_a}{\alpha}\log\frac{1}{p} + \sqrt{\frac{U_b}{\alpha_b}})$. Step 1.a and 1.b take time at most $O(\log n)$. Step 1.c takes time $O(\sigma\log n) = O(\frac{1}{\alpha}\log\frac{1}{p}\log n))$ using the sample operation of $\mathcal{E}$ by Theorem 3.3 if the algorithm does not restart. Step 1.d takes time $O(\frac{U_a}{\alpha})$ using the list operation of $\mathcal{E}$ by Theorem 3.3. Step 2 takes time $O(\sqrt{|E^s|}) = O(\sqrt{\frac{U_b}{\alpha_b}})$. So the running time in Algorithm 7.1 is at most $O(\frac{\log n}{\alpha}\log\frac{1}{p}+\frac{U_a}{\alpha}+\sqrt{\frac{U_b}{\alpha_b}}) = O(\frac{U_a}{\alpha}\log\frac{1}{p}+\sqrt{\frac{U_b}{\alpha_b}})$. □

Claim 8.11. After each edge deletion, $\mathcal{A}$ reports failure or updates $\mathcal{E}$ and $\mathcal{E}^{2d}$ such that the invariant is maintained.

Proof. Suppose that $\mathcal{A}$ does not report failure. We will prove that the invariant is maintained.

F^d and F are forests in G^d and G, respectively, because after preprocessing they are forest, and after that we only add replacement edges into them. Moreover, $F \supseteq F^d$ by Fact 8.9. It remains to prove that F^d and F span G^d and G respectively. That is, we prove that after each deletion if a replacement edge exists, then Algorithm 8.2 can find one.

We first prove this for F^d. Suppose that $e = (u,v) \in F^d$ is deleted. Let $T_u^d \cup (u,v) \cup T_v^d$ be a connected component in F^d. It easy to see that $\partial_{G^d}(V(T^d))$ is the set of replacement edges in E^d. Suppose Algorithm 8.2 cannot find a replacement edge in E^d. This happens only in Step 1.d as we assume $\mathcal{A}$ does not report failure. In this case, all edges incident to T_u^d are listed. So a replacement edge really do not exist.

Next, we prove the case of F. Suppose that $e = (u,v) \in F$ is deleted. We also denote $T_u \cup (u,v) \cup T_v$ a connected component in F. It is easy to see that $\partial_{G^s}(V(T))$ is the set of replacement edges in E^s. In Step 2, find_edge(T_u) returns an edge iff $\partial_{G^s}(V(T)) \neq \emptyset$. So if we cannot get any edge from both $\partial_{G^d}(V(T^d))$ and $\partial_{G^s}(V(T))$, then a replacement edge really do not exist. □

Claim 8.12. If $\phi(G_b) \geq \alpha_b$, then Algorithm 8.2 reports failure in Step 1.c with probability at most $O(p)$.

Proof. We use the same argument as in Claim 8.6. Let C be the component of G^d containing T^d. As $\phi(G_b) \geq \alpha_b$, by Theorem 6.1 $\phi(C) \geq \alpha$ with probability $1 - O(p)$. Now, we assume $\phi(C) \geq \alpha$ and prove that Algorithm 8.2 fail with probability at most p.

By Step 1.b, $|V(T^d)| \leq |V(C)|/2$, so $\delta_C(V(T)) \geq \alpha|V(T)| - U_a$ because there are at most U_a deletions. As $|V(T)| \geq 2U_a/\alpha$, we have $\delta_C(V(T)) \geq \frac{\alpha}{2}|V(T)|$. Therefore, all σ sampled edges are not cut edge with probability $(1 - \frac{\delta_C(V(T))}{vol_C(V(T))}) \leq (1 - \frac{\alpha/2}{3})^\sigma \leq p$. □

We conclude with the following lemma:

Lemma 8.13. *After each edge deletion, $\mathcal{A}$ takes time $O(\frac{U_a}{\alpha}\log\frac{1}{p} + \sqrt{\frac{U_b}{\alpha_b}}) = O(n^{1/2-\epsilon}\log\frac{1}{p})$ and either*

1. *reports failure, or*
2. *outputs the list of edges to be added or removed from F to obtain new one spanning forest of G. Moreover, the underlying graphs and forested of the given augmented ET tree $\mathcal{E}$ and 2-dim ET tree $\mathcal{E}^{2d}$ are updated such that $(G_\mathcal{E}, F_\mathcal{E}) = (G^d, F^d)$ where F^d is a spanning forest of G^d, and $(G_{\mathcal{E}^{2d}}, F_{\mathcal{E}^{2d}}) = (E^s \cup F^d, F)$ where F is a spanning forest of G and $F \supseteq F^d$.*

Moreover, if $\phi(G_b) \geq \alpha_b$, $\mathcal{A}$ does not report failure with probability $1 - O(p)$ for each update.

Combining Lemma 8.7 and Lemma 8.13, we immediately obtain Lemma 8.3. Moreover, Lemma 8.8 justifies Remark 8.1 when $D = \emptyset$.

8.2 Reduction to Broken Expanders

In this section, we prove Lemma 8.2 by using the algorithm from Lemma 8.3 as a main subroutine in a standard technique. Let $\mathcal{A}$ be the algorithm for Lemma 8.3. We would like to devise an algorithm $\mathcal{B}$ for Lemma 8.2.

Let G_b be the initial graph which is given to $\mathcal{B}$ as inputs. $\mathcal{B}$ preprocess G_b as described in Algorithm 8.3. We use two instances of augmented ET trees: $\mathcal{E}_1$ and $\mathcal{E}_2$, two instances of 2-dim ET trees: $\mathcal{E}_1^{2d}$ and $\mathcal{E}_2^{2d}$, and two instances of the algorithm $\mathcal{A}$ for Lemma 8.3: $\mathcal{A}_1$ and $\mathcal{A}_2$. We will maintain three spanning forests F_1 and F_2 and F_{final}.

1. Find a spanning tree F_b of G_b.
2. Set $F_1, F_2, F_{final} = F_b$.
3. Preprocess (G_b, F_b) using two instances of augmented ET trees $\mathcal{E}_1$ and $\mathcal{E}_2$.
4. Preprocess (F_b, F_b) using two instances of 2-dim ET trees $\mathcal{E}_1^{2d}$ and $\mathcal{E}_2^{2d}$.
5. Give $(n, p, G_b, \alpha_b, D, F_1, \mathcal{E}_1, \mathcal{E}_1^{2d})$ to $\mathcal{A}_1$ as input where $D = \emptyset$.

Algorithm 8.3: The algorithm for preprocessing the initial graph $G_b = (V, E)$

Claim 8.14. $\mathcal{B}$ takes $O(n)$ time to preprocess G_b and does not report failure. This running time is independent from the length of the update sequence.

Proof. Consider Algorithm 8.3. Step 1 takes $O(n)$ time. Step 2 is trivial. Step 3 and 4 take $O(n)$ time by Theorem 3.3 and Theorem 4.2. Step 6 takes no times by Remark 8.1. □

Throughout the update sequence, only the updates of the spanning forest F_{final} will be returned as the answer of $\mathcal{B}$. F_1 and F_2 are maintained by $\mathcal{A}_1$ and $\mathcal{A}_2$ respectively, and they are only for helping us maintaining F_{final}. $\mathcal{B}$ will report failure whenever $\mathcal{A}_1$ or $\mathcal{A}_2$ report failure.

Suppose that the update sequence of $\mathcal{B}$ is of size at most U_b. We divides the sequence into phases of size $U_a/2$. We will first describe an algorithm how to maintain F_1 and F_2 such that the following holds: in odd (even) phase i, $\mathcal{A}_1$ ($\mathcal{A}_2$) is maintaining a spanning forest F_1 (F_2) of the graph of during the whole phase.

For each odd phase, suppose that the first update of the phase is the $(i_0 + 1)$-th update. We assume that $\mathcal{A}_1$ is maintaining F_1 of the current graph G_{i_0} and can handle more $U_a/2$ updates for the whole phase. We will describe how to prepare $\mathcal{A}_2$ so that, at the end of the phase, i.e. after the $(i_0 + \frac{U_a}{2})$-th update, $\mathcal{A}_2$ is maintaining F_2 which is a spanning forest of the current graph $G_{i_0+\frac{U_a}{2}}$ at that time, and can handle $\frac{U_a}{2}$ more updates for the whole next phase. Let $(G_{\mathcal{E}_2}, F_{\mathcal{E}_2})$ and $(G_{\mathcal{E}_2^{2d}}, F_{\mathcal{E}_2^{2d}})$ be the underlying graphs and forests of $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ respectively.

For time period $[i_0+1, i_0+\frac{U_a}{4}]$, we distribute evenly the time needed for (i) setting $(G_{\mathcal{E}_2}, F_{\mathcal{E}_2}) = (G_{i_0}, F_{final,i_0})$ and $(G_{\mathcal{E}_2^{2d}}, F_{\mathcal{E}_2^{2d}}) = (F_{final,i_0}, F_{final,i_0})$ where G_{i_0} and F_{final,i_0} are an input graph G and the maintained forest F_{final} after the i_0-th update, respectively, and (ii) $\mathcal{A}_2$ to preprocess $(n, p, G_b, \alpha_b, D_{i_0}, F_{final,i_0}, \mathcal{E}_2, \mathcal{E}_2^{2d})$ where D_{i_0} is the set of deleted edge into time i_0 (Note that this is a valid input for $\mathcal{A}_2$).

Claim 8.15. During $[i_0+1, i_0+\frac{U_a}{4}]$, the procedure as described above takes $O(((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma+\frac{U_b}{\alpha_b^{1+1/\epsilon'}})\log\frac{1}{p})/U_a)$ time per update.

Proof. First of all, by Lemma 8.3, $\mathcal{A}_2$ preprocesses in time $P = O((\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}})\log\frac{1}{p})$ given a valid input. We just need to show that we can set $(G_{\mathcal{E}_2}, F_{\mathcal{E}_2}) = (G_{i_0}, F_{final,i_0})$ and $(G_{\mathcal{E}_2^{2d}}, F_{\mathcal{E}_2^{2d}}) = (F_{final,i_0}, F_{final,i_0})$ in at most $O(P)$ as well.

One way to argue this is that, at time $i_0 - U_a$, we know $(G_{\mathcal{E}_2}, F_{\mathcal{E}_2}) = (G_{i_0-U_a}, F_{final,i_0-U_a})$ and $(G_{\mathcal{E}_2^{2d}}, F_{\mathcal{E}_2^{2d}}) = (F_{final,i_0-U_a}, F_{final,i_0-U_a})$ by induction. Also, from time $i_0 - U_a$ to i_0, we have spent time on $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ at most $O(P)$ time. By remembering the execution log, we can revert both $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ to their states at time $i_0 - U_a$ in time $O(P)$. Now, as $G_{i_0-U_a}$ and F_{final,i_0-U_a} differ from G_{i_0} and F_{final,i_0} by at most $O(U_a)$ edges, we can update both $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ in time $O(U_a) \times \tilde{O}(\sqrt{\frac{U_b}{\alpha_b}}) = o(P)$[14]. As there are $U_a/4$ steps to distribute the work time, the process takes time per update as desired. □

For time period $[i_0 + \frac{U_a}{4} + 1, i_0 + \frac{U_a}{2}]$, at each time $i_0 + \frac{U_a}{4} + k$ for $1 \le k \le \frac{U_a}{4}$, we feed the $i_0 + (2k-1)$-th update and the $(i_0 + 2k)$-th update to $\mathcal{A}_2$. This takes $O(\frac{U_a}{\alpha}\log\frac{1}{p} + \sqrt{\frac{U_b}{\alpha_b}}) = O(n^{1/2-\epsilon}\log\frac{1}{p})$ time per update. Therefore, at the end of the phase, i.e. after the $(i_0 + \frac{U_a}{2})$-th update, $\mathcal{A}_2$ is maintaining F_2 which is a spanning tree of the current graph $G_{i_0+\frac{U_a}{2}}$, and can handle $\frac{U_a}{2}$ more updates as desired. In even phases, we do symmetrically for $\mathcal{E}_1$, $\mathcal{E}_1^{2d}$ and $\mathcal{A}_1$.

In the above description, we need to maintain F_{final} altogether with F_1 and F_2 because at the beginning of each phase, we set $(G_{\mathcal{E}_2}, F_{\mathcal{E}_2}) = (G_{i_0}, F_{final,i_0})$ and $(G_{\mathcal{E}_2^{2d}}, F_{\mathcal{E}_2^{2d}}) = (F_{final,i_0}, F_{final,i_0})$. To maintain F_{final} given that we know that F_1 and F_2 are correct spanning forests in odd and even phase respectively, we can use the exactly same idea as in Remark B.1 using additional $O(\log^2 n)$ update time.

Claim 8.16. For each update, $\mathcal{B}$ takes time at most $O(n^{1/2-\epsilon}\gamma\log\frac{1}{p})$ when $U_b = n^{1/2+\epsilon}$, $\alpha_b = 1/n^{2\epsilon}$, $U_a = n^{1/2-\epsilon}\alpha_b^{1/\epsilon'} = n^{1/2-\epsilon-2\epsilon/\epsilon'}$.

Proof. From the above discussion, we have that the update time of $\mathcal{B}$ is

$$
\begin{aligned}
& O(n^{1/2-\epsilon}\log\frac{1}{p} + \frac{(\frac{{U_b}^{1.5+\epsilon'}}{\alpha_b^{4+\epsilon'}}\gamma + \frac{U_b}{\alpha_b^{1+1/\epsilon'}})\log\frac{1}{p}}{U_a} + \log^2 n) \\
= \;& O(\gamma\log\frac{1}{p}) \times O(n^{1/2-\epsilon} + \frac{n^{(1.5+\epsilon')(1/2+\epsilon)+2\epsilon(4+\epsilon')} + n^{1/2+\epsilon+2\epsilon(1+1/\epsilon')}}{n^{1/2-\epsilon-2\epsilon/\epsilon'}}) \\
= \;& O(\gamma\log\frac{1}{p}) \times O(n^{1/2-\epsilon} + \frac{n^{(1.5+\epsilon')(1/2+\epsilon)+2\epsilon(4+\epsilon')}}{n^{1/2-\epsilon-2\epsilon/\epsilon'}}) \\
= \;& O(\gamma\log\frac{1}{p}) \times O(n^{1/2-\epsilon} + \frac{n^{3/4+9.5\epsilon+3\epsilon\epsilon'+\epsilon'/2}}{n^{1/2-\epsilon-2\epsilon/\epsilon'}}) \\
= \;& O(\gamma\log\frac{1}{p}) \times O(n^{1/2-\epsilon} + n^{1/4+10.5\epsilon+3\epsilon\epsilon'+2\epsilon/\epsilon'+\epsilon'/2}) \\
= \;& O(n^{1/2-\epsilon}\gamma\log\frac{1}{p})
\end{aligned}
$$

[14] Another way to argue this is by making $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ *fully persistent*, we can just "go back in time" to start updating $\mathcal{E}_2$ and $\mathcal{E}_2^{2d}$ from time $i_0 - U_a$. The overhead factor for making it fully persistent is $O(\log\log n)$ by Straka [37] (cf. [10, 22]).

when $\epsilon = 0.0069459$ and $\epsilon' = \sqrt{\frac{2\epsilon}{3\epsilon+0.5}} = 0.16332$. □

Claim 8.17. If $\phi(G_b) \geq \alpha_b$, then, for each update, $\mathcal{B}$ does not report failure with probability $1-O(p)$.

Proof. If $\phi(G_b) \geq \alpha_b$, then, for each update, both $\mathcal{A}_1$ and $\mathcal{A}_2$ do not fail with probability $1 - O(p)$ by Lemma 8.3. As $\mathcal{B}$ only reports failure when $\mathcal{A}_1$ or $\mathcal{A}_2$ does, we are done. □

The resulting algorithm $\mathcal{B}$ from the above reduction is the desired algorithm for Lemma 8.2.

Lemma 8.18 (Restatement of Lemma 8.2)**.** *For any n, p and T, there is a decremental* SF *algorithm $\mathcal{B}$ for 3-bounded degree graphs with n nodes that works correctly against adaptive adversaries with certainty for the first T updates. $\mathcal{B}$ preprocesses an initial graph G_b in time $t_p(n,p) = O(n)$. Given each edge deletion, $\mathcal{B}$ takes $t_u(n,p,T)$ worst-case time to update or report failure, where $t_u(n,p,n^{1/2+\epsilon}) = O(n^{1/2-\epsilon}\gamma \log \frac{1}{p})$. If $\phi(G_b) \geq 1/n^{2\epsilon}$, then, for each update, $\mathcal{B}$ does not report failure with probability $1 - O(p)$.*

Proof of Lemma 8.2. By Claim 8.14, $\mathcal{B}$ has preprocessing time $t_p(n,p) = O(n)$ independent of T. Given each edge deletion, $\mathcal{B}$ takes $t_u(n,p,T)$ worst-case time to update or report failure where $t_u(n,p,n^{1/2+\epsilon}) = O(n^{1/2-\epsilon}\gamma \log \frac{1}{p})$ by Claim 8.16. Finally, by Lemma 8.3, if $\phi(G_b) \geq \alpha_b = 1/n^{2\epsilon}$, then, for each update, $\mathcal{B}$ does not report failure with probability $1 - O(p)$. □

References

[1] Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinninger, and Richard Peng. On fully dynamic graph sparsifiers. *CoRR*, abs/1604.02094, 2016. 22

[2] Stephen Alstrup, Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Transactions on Algorithms*, 1(2):243–264, 2005. 13, 15, 54

[3] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In *ITCS*, pages 311–319. ACM, 2016. 5, 24

[4] Shai Ben-David, Allan Borodin, Richard M. Karp, Gábor Tardos, and Avi Wigderson. On the power of randomization in online algorithms. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 379–386. ACM, 1990. 2

[5] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 798–805, Sept 2008. 3, 6, 17, 56, 57

[6] Aaron Bernstein and Shiri Chechik. Deterministic decremental single source shortest paths: beyond the o(mn) bound. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 389–397, 2016. 2

[7] Aaron Bernstein and Shiri Chechik. Deterministic partially dynamic single source shortest paths for sparse graphs. In *SODA*, 2017. 2

[8] Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. Fully dynamic maximum matching and vertex cover in $o(\log^3 n)$ worst case update time. *Manuscript*, 2016. 1

[9] Michael R. Capalbo, Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 659–668, 2002. 17, 56

[10] James R. Driscoll, Neil Sarnak, Daniel Dominic Sleator, and Robert Endre Tarjan. Making data structures persistent. *J. Comput. Syst. Sci.*, 38(1):86–124, 1989. Announced at STOC 1986. 47

[11] David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification - a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, 1997. Announced at FOCS 1992. 1, 3, 4, 20, 21, 62

[12] Greg N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM J. Comput.*, 14(4):781–798, 1985. Announced at STOC'83. 1, 3, 4, 13, 20, 21

[13] David Gibb, Bruce M. Kapron, Valerie King, and Nolan Thorn. Dynamic graph connectivity with improved worst case update time and sublinear space. *CoRR*, abs/1509.06464, 2015. 1, 2, 3, 6, 19

[14] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *55th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 18-21, 2014, Philadelphia, PA, USA*, pages 146–155, 2014. 2

[15] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the o(mn) barrier and derandomization. *SIAM J. Comput.*, 45(3):947–1006, 2016. Announced at FOCS'13. 2

[16] Monika Rauch Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. ACM*, 46(4):502–516, 1999. Announced at STOC 1995. 1, 2, 3, 4, 6, 7, 11, 15, 17, 20, 21, 54, 57, 58, 59, 63

[17] Monika Rauch Henzinger and Valerie King. Maintaining minimum spanning forests in dynamic graphs. *SIAM J. Comput.*, 31(2):364–374, 2001. 1

[18] Monika Rauch Henzinger and Mikkel Thorup. Sampling to provide or to bound: With applications to fully dynamic graph algorithms. *Random Struct. Algorithms*, 11(4):369–379, 1997. 1

[19] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001. Announced at STOC 1998. 1, 20

[20] Shang-En Huang, Dawei Huang, Tsvi Kopelowitz, and Seth Pettie. Fully dynamic connectivity in $o(\log n(\log\log n)^2)$ amortized expected time. In *SODA*, 2017. 1

[21] Donggu Kang and James Payor. Flow rounding. *CoRR*, abs/1507.08139, 2015. 65

[22] Haim Kaplan. Persistent data structures. In *In handbook on data structures and applications*, 1995. 47

[23] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1131–1142. SIAM, 2013. 1, 2, 3, 6, 17, 19, 20, 70

[24] Casper Kejlberg-Rasmussen, Tsvi Kopelowitz, Seth Pettie, and Mikkel Thorup. Faster worst case deterministic dynamic connectivity. In Piotr Sankowski and Christos D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPIcs*, pages 53:1–53:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. 2, 3, 4, 20, 21

[25] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In Chandra Chekuri, editor, *SODA*, pages 217–226. SIAM, 2014. 3, 5, 23

[26] Rohit Khandekar, Satish Rao, and Umesh V. Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4), 2009. 3, 5, 23, 24, 64, 65, 66

[27] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999. 5, 23

[28] Aleksander Madry. Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 121–130, 2010. 2

[29] Sotiris E. Nikoletseas, John H. Reif, Paul G. Spirakis, and Moti Yung. Stochastic graphs have short memory: Fully dynamic connectivity in poly-log expected time. In *Automata, Languages and Programming, 22nd International Colloquium, ICALP95, Szeged, Hungary, July 10-14, 1995, Proceedings*, pages 159–170, 1995. 12

[30] Lorenzo Orecchia and Zeyuan Allen Zhu. Flow-based algorithms for local graph clustering. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1267–1286, 2014. 3, 9, 29, 67, 68, 69

[31] Mihai Patrascu and Erik D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM J. Comput.*, 35(4):932–963, 2006. Announced at SODA'04 and STOC'04. 1

[32] Mihai Patrascu and Mikkel Thorup. Planning for fast connectivity updates. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 263–271, 2007. 1, 4

[33] Richard Peng. Approximate undirected maximum flows in $O(m\text{polylog}(n))$ time. In *SODA*, pages 1862–1867. SIAM, 2016. 3, 5, 23, 66

[34] Liam Roditty and Uri Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 499–508. IEEE Computer Society, 2004. 2

[35] Jonah Sherman. Nearly maximum flows in nearly linear time. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 263–269, 2013. 3, 5, 23

[36] Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983. 11, 54, 65

[37] Milan Straka. *Functional Data Structures and Algorithms.* PhD thesis, Charles University, 2013. 47

[38] Mikkel Thorup. Near-optimal fully-dynamic graph connectivity. In F. Frances Yao and Eugene M. Luks, editors, *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 343–350. ACM, 2000. 1, 2

[39] Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 112–119. ACM, 2005. 1

[40] Mikkel Thorup. Fully-dynamic min-cut. *Combinatorica*, 27(1):91–127, 2007. Announced at STOC'01. 3, 13, 15

[41] Christian Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1757–1769, 2013. 1

A Formalization of Dynamic Algorithms vs. Adversaries

In this section, we formalize the notions of dynamic algorithms and adversaries themselves. This can be of independent interest because, to be best of our knowledge, these notions were not defined formally in the literature in the context of dynamic algorithms yet.[15]

First, we define *dynamic problems.* A dynamic problem $\mathcal{P}$ is specified by a tuple $(\mathcal{G}, U, Q, A, \mathcal{S}, Corr)$ with the following definitions.

- $\mathcal{G}, U, Q$ and A are the sets of valid *underlying objects*, *updates*, *queries*, and *answers*, respectively. Each update $u \in U$ is a function $u : \mathcal{G} \to \mathcal{G}$. We call $U \cup Q$ a set of valid *operations.*
- Let $\mathcal{S} = \bigcup_{i \geq 0} \mathcal{S}_i$ where, for any i, $\mathcal{S}_i \subseteq \mathcal{G} \times (U \cup Q)^i$ is a set of valid *operation sequences* of length i. For any $S_i = (G_0, o_1, \ldots, o_i) \in \mathcal{S}_i$, G_0 is called an *initial* object. For any $1 \leq j \leq i$, we say G_j is the updated underlying object after the j-th operation. More precisely, G_j is obtains by iteratively applying to G_0 each update operation o_k for $k = 1, \ldots, j$ where $o_k \in U$.
- Given an operation sequence $S_i = (G_0, o_1, \ldots, o_i) \in \mathcal{S}_i$ and an answer sequence $\vec{a}_i = (a_0, a_1, \ldots, a_i) \in A^{i+1}$, we say $\pi_i = (G_0, a_0, o_1, a_1, \ldots, o_i, a_i)$ is a *transcript* at round i. We sometimes write $\pi_i = (S_i, \vec{a}_i)$. Let Π_i be the set of all transcripts at round i, and $\Pi = \bigcup_{i \geq 0} \Pi_i$.
- For any $(\pi_i, o_{i+1}) \in \Pi_i \times (U \cup Q)$, $Corr(\pi_i, o_{i+1})$ is a set of *correct answers* for o_{i+1} given π_i. If $o_{i+1} \in Q$ is a query, then $Corr(\pi_i, o_{i+1}) \subseteq A$. If $o_{i+1} \in U$ is an update, then for convenience we define $Corr(\pi_i, o_i) = \{\emptyset\}$.

To illustrate, we give some examples how the tuple captures the definitions of many dynamic problems in the literature.

Example A.1 (Dynamic Connectivity)**.** *$\mathcal{G}$ is the set of all undirected graphs. Each update $u \in U$ specifies which edge in the graph to be inserted or deleted. Each query $q \in Q$ specifies a pair of nodes. The set of answers $A = \{YES, NO\}$. Given a transcript $\pi_i \in \Pi_i$ and a query $(u, v) \in Q$, $Corr(\pi_i, (u, v)) = \{YES\}$ if u and v is connected in G_i and $Corr(\pi_i, (u, v)) = \{NO\}$ otherwise. For any valid sequence $S \in \mathcal{S}$, S always excludes trivially invalid updates and queries (e.g. a deletion of an edge e when e is not in the graph, or a query (u, v) when a node u is not in the graph).*

Example A.2 (Dynamic Spanning Tree)**.** *Similarly, we have $\mathcal{G}$ is the set of all undirected graphs. Each update $u \in U$ specifies which edge in the graph to be inserted or deleted. The set of query is trivial: $Q = \{\mathsf{list}\}$. For any sequence $S_i = (G_0, o_1, \ldots, o_i) \in \mathcal{S}_i$, we have $o_j \in U$ for odd j and $o_j \in Q$ for even j, so that this captures the fact that we immediately query after each update in this problem[16]. The set of answers A is the list of edges to be added or removed from some set. So given any sequence $a_0, \ldots, a_i \in A^{i+1}$, this defines a set F_i of edges (by first applying a_0 to an empty set, then applying a_1, and so on). Let $\pi_i = (G_0, a_0, \ldots, o_i, a_i) \in \Pi_i$, we have $a_{i+1} \in Corr(\pi_i, \mathsf{list})$ iff F_{i+1} is a spanning forest of G_{i+1} where F_{i+1} is obtained by applying a_{i+1} to F_i.*

Example A.3. *For the dynamic problems on planar graphs, $\mathcal{G}$ is the set of all planar graphs, and $\mathcal{S}$ is defined in such a way that after each update the underlying graph remains planar. We can similar define this for the problem on bounded arboricity graphs.*

[15] The notion of adversaries has been defined formally on the context of online algorithms whose focus is on competitiveness [?, ?], but has not been formalized for dynamic algorithms whose focuses are on update and query time.

[16] In \cite[Trade-off-MST], they define the variant of the problem where a query need not immediately follows each update, and show the application of this variant.

Next, we can define the notion of adversaries and algorithms for dynamic problems. Fix some dynamic problem $\mathcal{P}$. An *adversary* f is a function $f : \Pi \to (U \cup Q)$ (except $f(\cdot) = G_0 \in \mathcal{G}$) such that for any $\pi_i = (S_i, \vec{a}_i) \in \Pi_i$ we have $(S_i, f(\pi_i)) \in \mathcal{S}_{i+1}$. For all $(S, \vec{a}), (S, \vec{a}') \in \Pi$ where $\vec{a} \neq \vec{a}'$, if we have $f(S, \vec{a}) = f(S, \vec{a}')$, then we say that f is *oblivious*, otherwise f is *adaptive* in general.

If an adversaries f is oblivious, then f is a function of this form $f : \mathcal{S} \to (U \cup Q)$. For any oblivious f, f can be written as a function of the form $f : \mathcal{S} \to (U \cup Q)$ and the *operation sequence generated by* f is $S_f = (G_0, o_1, o_2, \dots)$ where $G_0 = f(\cdot)$ and $o_i = f(G_0, \dots, o_{i-1})$ for all $i \geq 1$.

An *algorithm* $\mathcal{A}$ is a function $\mathcal{A} : \{0,1\}^\infty \times (\mathcal{G} \cup (\Pi \times (U \cup Q))) \to A \cup \{\emptyset\}$ such that for any infinite string $R \in \{0,1\}^\infty$ and transcript $\pi \in \Pi$, we have $\mathcal{A}(R, \pi, o) = \emptyset$ if $o \in U$, and $\mathcal{A}(R, \pi, o) \in A$ if $o \in Q$.

Given any adversary f, algorithm $\mathcal{A}$ and an infinite string R, this determines a sequence $(G_0, a_0, o_1, a_1, o_2, a_2, \dots)$ where $G_0 = f(\cdot)$, $a_0 = \mathcal{A}(R, G_0)$, $o_i = f(G_0, a_0, \dots, o_{i-1}, a_{i-1})$, and $a_i = \mathcal{A}(R, G_0, a_0, \dots, o_{i-1}, a_{i-1}, o_i)$ for $i \geq 1$.[17] For $i \geq 0$, let $\pi_i = (G_0, a_0, \dots, o_i, a_i) \in \Pi_i$ denote a transcript at round i given $f, \mathcal{A}$ and R. We can write $o_i = f(\pi_{i-1})$ and $a_i = \mathcal{A}(R, \pi_{i-1}, o_i)$. Let t_i be the time period $\mathcal{A}$ uses for outputting a_i starting counting from the time a_{i-1} is outputted. We emphasize that all variables in this paragraph are determined once $f, \mathcal{A}$ and R are fixed. Let $succ_i(f, \mathcal{A}, R)$ be the successful event that occurs when $a_i \in Corr(\pi_{i-1}, o_i)$ (or $a_0 \in Corr(G_0)$ when $i = 0$).

Now, we are ready to define correctness of an algorithm.

Definition A.4 (Correctness)**.** *Let $\mathcal{F}$ be some set of adversaries. An algorithm $\mathcal{A}$ works correctly against $\mathcal{F}$ with probability p for the first t steps, if, for every $f \in \mathcal{F}$, $\Pr_R[succ_i(f, \mathcal{A}, R)] \geq p$ for all $0 \leq i \leq t$. If $t = \infty$, then the phrase "for the first t steps" is omitted.*

Next, we define the precise meaning of preprocessing, update and query time of an algorithm.

Definition A.5 (Running time)**.** *Let $\mathcal{F}$ be some set of adversaries. An algorithm $\mathcal{A}$ has worst-case preprocessing time T_p, worst-case update time T_u and worst-case query time T_q against $\mathcal{F}$ with probability p if, for every for every $f \in \mathcal{F}$, we have $\Pr_R[t_0 \leq T_p] \geq p$ and, for each $i \geq 1$, $\Pr_R[t_i \leq T_u] \geq p$ when $o_i \in U$ and $\Pr_R[t_i \leq T_q] \geq p$ when $o_i \in Q$.*

If $\mathcal{F}$ is a set of all adversaries, then we say that $\mathcal{A}$ works *against adaptive adversaries*. If $\mathcal{F}$ is a set of all oblivious adversaries, then we say that $\mathcal{A}$ works *against oblivious adversaries*.

From above definition, adversaries are deterministic. We can also define a randomized version, although as will be shown below, this does not give more power. A *randomized adversary* f is a function $f : \{0,1\}^\infty \times \Pi \to (U \cup Q)$ where for any string R', $f_{R'} = f(R', \cdot)$ is an adversary. We say that f is an *oblivious randomized adversary* if, for all R', $f_{R'}$ is an oblivious adversary.

Definition A.6 (Correctness (randomized adversaries))**.** *Let $\mathcal{F}$ be some set of randomized adversaries. An algorithm $\mathcal{A}$ works correctly against $\mathcal{F}$ with probability p for the first t steps, if, for every $f \in \mathcal{F}$, $\Pr_{R,R'}[succ_i(f_{R'}, \mathcal{A}, R)] \geq p$ for all $0 \leq i \leq t$. If $t = \infty$, then the phrase "for the first t steps" is omitted.*

The following proposition shows that, to design an algorithm that works against adaptive adversaries or oblivious adversaries, it is safe to assume that the adversaries are deterministic.

Proposition A.7. *Let $\mathcal{F}$ and $\mathcal{F}_r$ and be a set of all (oblivious) deterministic adversaries and (oblivious) randomized adversaries respectively. If $\mathcal{A}$ works correctly against $\mathcal{F}$ with probability p for the first t steps, then so does against $\mathcal{F}_r$.*

[17] We note the resemblance of this definitions and the transcript in the context of communication complexity is intentional.

Proof. Let $f \in \mathcal{F}_r$. We have that for all $0 \le i \le t$, $\Pr_{R,R'}[succ_i(f_{R'}, \mathcal{A}, R)] \ge \min_{R'} \Pr_R[succ_i(f_{R'}, \mathcal{A}, R)] \ge p$ where the last inequality is because, for any R', we have $f_{R'} \in \mathcal{F}$. □

To say instead that $\mathcal{A}$ has *amortized* or *expected worst-case* or *expected amortized* update time T_u, the expression $t_i \le T_u$ is replaced by $\sum_{j=1}^{i} t_i \le i \times T_u$ or $E_R[t_i] \le T_u$ or $E_R[\sum_{j=1}^{i} t_i] \le i \times T_u$, respectively. Similarly, for query time and preprocessing time. A Monte Carlo algorithm works correctly with some probability and has guarantee on running time without expectation. A Las Vegas algorithm work correctly with certainty but its running time is guaranteed only in expectation. A deterministic algorithm both works correctly with certainty and has guarantee on running time without expectation.

In many dynamic problems including dynamic SF problem, each update is immediately followed by a query. There are many other examples including dynamic MSF, and most, if not all, dynamic problems which maintains an approximate/exact value (e.g. the size of maximum matching, the size of a global min cut, etc.) For these dynamic problems, if the update and query time from the definition above are T_u and T_q, then the update time from the original definition is $T_u + T_q$.

B Omitted proof

B.1 Proof of Theorem 3.3

Proof. By using *ET tree* data structure from [16, Section 2.1.4 - 2.1.5], we can preprocess G and F in time $O(m)$ and handle all operations (except path) in time $O(\log n)$, and handle tree_size and count_tree in $O(1)$ time.

In order to also handle path, it suffices to show that there is a data structure that preprocess G and F in time $O(m)$ and handle insert_tree, delete_tree and path in time $O(\log n)$. This is because we can just implement this data structure in parallel with ET tree. A data structure called *link-cut tree* from [36] or *top tree* from [2] can indeed do the job. From [2, Section 2], top tree can handle insert_tree, delete_tree. In order to handle path(u, v, i), we can the operation expose(u, v) (defined in [2]) which allows us to do binary search on the unique path P_{uv} from u to v in F and find the i-th vertex in P_{uv}. □

B.2 From Section 4.1 and Section 4.2

B.2.1 Proof of Theorem 4.5

To describe the algorithm for Theorem 4.5, we need to define several notions: *base clusters*, *non-base clusters*, *clusters*, and *hierarchies of clusters*.

For any list $L = (e_1, e_2, \ldots, e_s) \in \mathcal{L}$, we call a collection of sublists $B_1, \ldots, B_g$ of L as *base clusters* of L if the following hold: 1) $g = O(\sqrt{k})$, 2) $L = B_1 \ldots B_g$, and 3) $\min\{\sqrt{k}/2, |L|\} \le |B_i| \le 2\sqrt{k}$. That is, the base clusters partition L into $O(|L|/\sqrt{k})$ parts where each part of length around $\sqrt{k}$ (or as long as $|L|$). Let $\mathcal{B}(L)$ denote the set of base clusters of L. Let $\mathcal{B}(\mathcal{L}) = \bigcup_{L \in \mathcal{L}} \mathcal{B}(L)$ denote the set of base clusters of $\mathcal{L}$.

Let $\mathcal{B}(L)$ be the base clusters of L. Let $\mathcal{T}_L$ be a balanced binary search tree whose leaves correspond to the base clusters $\mathcal{B}(L)$. Each node $u \in \mathcal{T}_L$ corresponds to a sublist C_u of L. If u is the i-th leaf, then $C_u = B_i$ is the i-th base cluster of L. If u is a non-leaf and have left child v and right child w, then $C_u = C_v C_w$ is the concatenation of the two sublists corresponding the two nodes v and w. For each non-leaf node $u \in \mathcal{T}_L$, we call C_u a *non-base cluster* of L. Note that, for the root $r \in \mathcal{T}_L$, $C_r = L$. We call C_r the *root cluster* of L. We denote $\mathcal{N}(L)$ the set of all non-base clusters of L, and $\mathcal{N}(\mathcal{L}) = \bigcup_{L \in \mathcal{L}} \mathcal{N}(L)$ denote the set of non-base clusters of $\mathcal{L}$. Similarly,

we call $\mathcal{T}_L$ the *hierarchy of clusters* in L, and $\mathcal{T}_{\mathcal{L}} = \bigcup_{L\in\mathcal{L}} \mathcal{T}_L$ be the set of such hierarchies. Let $\mathcal{C}(L) = \mathcal{B}(L) \cup \mathcal{N}(L)$ denote the set of *clusters* of L, and similarly $\mathcal{C}(\mathcal{L}) = \bigcup_{L\in\mathcal{L}} \mathcal{C}(L)$ denote the set of clusters of $\mathcal{L}$.

Next, we describe the main data structure that for answering the intersection query. We say that A is an *intersection table of clusters in* $\mathcal{L}$ if for any two clusters $C_1, C_2 \in \mathcal{C}(\mathcal{L})$, $A(C_1, C_2) = e$ if a number $e \in C_1 \cap C_2$ exists, otherwise, $A(C_1, C_2) = \emptyset$.

Proposition B.1. *Let $\mathcal{L}$ be a collection of lists. Given an intersection table of clusters in $\mathcal{L}$, for any $L_1, L_2 \in \mathcal{L}$, a query* intersection(L_1, L_2) *can be answer in $O(1)$ time.*

Proof. Let C_1 and C_2 be the root clusters of L_1 and L_2. Recall that $C_1 = L_1$ and $C_2 = L_2$. We can just return $A(C_1, C_2)$ to the query intersection(L_1, L_2). □

Now, the goal is to prove the following. We will exploit the assumptions from Theorem 4.5 so that we can preprocess and update quickly.

Proposition B.2. *Let $\mathcal{L} = \{L_i\}_i$ be a collection of lists such that $\sum_i |L_i| \le k$ and the number of lists in $\mathcal{L}$ is bounded by some constant. There are at most $O(\sqrt{k})$ clusters of $\mathcal{L}$, i.e. $|\mathcal{C}(\mathcal{L})| = O(\sqrt{k})$.*

Proof. As there are only a constant number of lists in $\mathcal{L}$, we only any to bound $|\mathcal{C}(L)|$ for each $L \in \mathcal{L}$. By definition of base clusters, it is clear that $|\mathcal{B}(L)| = O(\sqrt{k})$ because $|L| \le k$, and so $|\mathcal{N}(L)| = O(\sqrt{k})$. Hence $|\mathcal{C}(L)| \le |\mathcal{B}(L)| + |\mathcal{N}(L)| = O(\sqrt{k})$. □

We note that $\mathcal{N}(\mathcal{L})$ and $\mathcal{C}(\mathcal{L})$ and is defined based on the base clusters $\mathcal{B}(\mathcal{L})$ and the hierarchies of clusters $\mathcal{T}_{\mathcal{L}}$. Therefore, in our data structure, we will only explicitly maintain $\mathcal{B}(\mathcal{L})$ and $\mathcal{T}_{\mathcal{L}}$ and the intersection table A of clusters in $\mathcal{L}$. We show how to build and maintain quickly the first two in Lemma B.3, and then do the same for the intersection table in Lemma B.4.

Lemma B.3. *Let $\mathcal{L}$ be a collection of lists such that the total length of lists in $\mathcal{L}$ is bounded by $\sum_{L\in\mathcal{L}} |L| \le k$. Then the base clusters $\mathcal{B}(\mathcal{L})$ and the hierarchies of cluster $\mathcal{T}_{\mathcal{L}}$ of $\mathcal{L}$ can be built in $O(k)$ time and can be maintained under* create, merge, *and* split *in $O(\log k)$ time.*

Proof. To build the base clusters $\mathcal{B}(\mathcal{L})$, we just read through $\mathcal{L}$ and split it into sublists of size $\sqrt{k}$. After having the base clusters $\mathcal{B}(\mathcal{L})$, for each $L \in \mathcal{L}$, we just build $\mathcal{T}_L$ as some balanced binary search tree whose leaves correspond $\mathcal{B}(L)$. The total time is obviously $O(k)$. Next, we show how to handle updates.

It is trivial to handle create. Given either merge or split to $\mathcal{L}$, to maintain $\mathcal{B}(\mathcal{L})$, observe that we only need to merge and split the base clusters by constant number of time. We will store each base cluster of $\mathcal{L}$ in a balanced binary search tree (binary search tree with logarithmic depth). As we can merge and split balanced binary search tree with at most k elements in time $O(\log k)$, we are done.

To maintain $\mathcal{T}_{\mathcal{L}}$, for each L, $\mathcal{T}_L$ is simply some balanced binary search tree whose leaves are base clusters of L. If some base cluster of L is splitted or merged, we maintain $\mathcal{T}_L$ as a balanced binary search tree in time $O(\log k)$. Because either merge, or split to $\mathcal{L}$ only generate a constant number of merging and splitting to base clusters, we are done. □

Lemma B.4. *Let $\mathcal{L} = \{L_i\}_i$ be a collection of lists. Suppose that 1) the total length of lists in $\mathcal{L}$ is bounded by $\sum_i |L_i| \le k$, 2) the number of lists in $\mathcal{L}$ is bounded by some constant, and 3) each number occurs in $\mathcal{L}$ at most a constant number of times. The intersection table of clusters in $\mathcal{L}$ can be built in time $O(k)$ and can be maintained under* create, merge, *and* split *in $O(\sqrt{k}\log k)$ time.*

Proof. By Proposition B.2, we know that there are only $O(\sqrt{k})$ clusters. By the third assumption, we can maintain the following data structure: given any number e, all the base clusters $B \in \mathcal{B}(\mathcal{L})$, such that e occurs in B, are returned in $O(1)$ time. We first describe how to handle the update A. We will describe how to build A later.

It is trivial to handle create. We only show how to handle merge, and split. Given the operations merge or split to $\mathcal{L}$, we only only need to merge and split the base clusters $\mathcal{B}(\mathcal{L})$ by constant number of times. It suffices to show that intersection table of clusters in $\mathcal{L}$ can be maintained in $O(\sqrt{k}\log k)$ time under under each merge or split of base clusters in $\mathcal{B}(\mathcal{L})$.

Consider any list $L \in \mathcal{L}$. Let $\mathcal{B}(L)$ be the base clusters of L and $\mathcal{T}_L$ be the hierarchy of clusters of L. Suppose that two base clusters $B_1, B_2 \in \mathcal{B}(L)$ are merged into $B' = B_1B_2$. Let $\mathcal{B}'(L)$ and $\mathcal{T}'_L$ be the updated base clusters and hierarchy of clusters of L after merging. By Lemma B.3, we can obtain $\mathcal{B}'(L)$ and $\mathcal{T}'_L$ in time $O(\log k)$. Let P the path in $\mathcal{T}'_L$ from root of to the leaf corresponding to the new base cluster B'. Observe that, any cluster of C_u corresponding to a node $u \in \mathcal{T}'_L \setminus P$ outside the path P is not affected by the merging. Therefore, to obtain an updated intersection table A of clusters, it is enough to compute $A(C_u, C)$ for all $u \in P$ and $v \in \mathcal{T}'_L$. We will do this bottom-up.

First, we compute $A(B', C_v)$ for all $v \in \mathcal{T}'_L$ as follows. For each number $e \in B'$, for each base clusters B'' that $B'' \ni e$, and for each clusters C that C is an ancestor of B'' in $\mathcal{T}'_L$, we set $A(B', C) = e$. Other entries are set to $\perp$ by default. This can be done in $O(\sqrt{k}\log k)$ times. Because there are $O(\sqrt{k})$ numbers in B'', for each number e we can list all base clusters B'' containing e in $O(1)$, and each such base cluster B'' has at most $O(\log k)$ ancestor because $\mathcal{T}'_L$ has $O(\log k)$ depth.

Next, we compute $A(C_u, C_v)$ for all $v \in \mathcal{T}'_L$ where $u \in P$ is a non-leaf in $\mathcal{T}'_L$ and has u_1 and u_2 has children. For each $v \in \mathcal{T}'_L$, we set $A(C_u, C_v) = A(C_{u_1}, C_v)$ if $A(C_{u_1}, C_v) \neq \perp$, otherwise $A(C_u, C_v) = A(C_{u_2}, C_v)$. This can be done in $O(\sqrt{k})$ because there are only $O(\sqrt{k})$ clusters. In total, the time need to compute $A(C_u, C_v)$ for all $v \in \mathcal{T}'_L$ and for all non-leaves $u \in P$ is $O(\sqrt{k}\log k)$ because $\mathcal{T}'_L$ has $O(\log k)$ depth.

The correctness of the intersection table is clear. To build A, use the same idea as the way we update, it is obvious how to do it in $O(k \log k)$ but we can save the $O(\log k)$ factor by building it "bottom-up". □

By Proposition B.1 and Lemma B.4, this concludes the proof of Theorem 4.5.

B.2.2 Proof of Theorem 4.14

To prove the theorem, we need to define unbalanced expanders.

Definition B.5. *Fix any k, d, ϵ. A (k, ϵ)-*unbalanced expander *is a bipartite simple graph $G = ((L, R), E)$ with left degree d (i.e. for every $v \in L$, $\deg(v) = d$) such that for any $X \subset L$ with $|X| \leq k$, the set of neighbors $N(X)$ of X has size $|N(X)| \geq (1 - \epsilon)d|X|$.*

There is an explicit construction of a (k, ϵ)-unbalanced expander by [9].

Theorem B.6 (Theorem 7.3 of [9] and Proposition 7 of [5])**.** *Fix any $n \geq k$ and $\epsilon > 0$. There is a (k, ϵ)-unbalanced expander $G = ((L, R), E)$ with left degree $\beta_1 = 2^{O(\log(\log n/\epsilon))^3}$, $|L| = n$ and $|R| = k\beta_1/\epsilon^{O(1)}$. Given $u \in L$ and $i \in [\beta]$, we can compute the i-th edge incident to u in $\beta_2 = O(poly(\log(n), 1/\epsilon))$ time.*

Next, we need some more definitions about matrices. The *Hamming code matrix* (or *bit test matrix*) $H^{(n)} \in \{0,1\}^{\lceil \log_2 n+1 \rceil \times n}$ is a matrix whose the i-th column has 1 followed by the binary representation of i. Let P be a $p \times n$ matrix and Q be a $q \times n$ matrix with rows $\{P_i\}_{0 \leq i < p}$ and

$\{Q_j\}_{0\le j<q}$ respectively. The *row tensor product* $P \otimes_R Q$ of P and Q is a $pq \times n$ matrix whose rows are $\{P_i Q_j\}_{0\le i<p, 0\le j<q}$ where $P_i Q_j$ denotes the component-wise product of the two row vectors. The construction of k-sparse recovery linear sketch in [5] is as follows.

Theorem B.7 (Theorem 18 of [5]). *Fix and $n \ge k$, r and $\epsilon = 1/8$. Let $\Psi \in \{0,1\}^{r\times n}$ be a bi-adjacency matrix of a (k,ϵ)-unbalanced expander. Let $\Phi = \Psi \otimes_R H^{(n)}$. Then there is an algorithm $\mathcal{A}_{recover}$ such that, for any k-sparse vector v $v \in [-M, M]^n$ where $M = poly(n)$, given Φv, all non-zero entries of v can be identified in time $O(r \log^2 n)$.*

Now, we are ready to prove Theorem 4.14.

Proof of Theorem 4.14. By Theorem B.6 when setting $\epsilon = 1/8$, there is a bi-adjacency matrix $\Psi \in \{0,1\}^{O(k\beta_1)\times n}$ of a (k,ϵ)-unbalanced expander $G = ((L,R),E)$ with left degree $\beta_1 = 2^{O(\log\log n)^3}$. Then, by Theorem B.7, we have that $\Phi = \Psi \otimes_R H^{(n)} \in \{0,1\}^{d\times n}$ where $d = O(k\beta_1 \log n)$ and there is algorithm $\mathcal{A}_{recover}$ for recovering k-sparse vector v, given Φv, in time $O(k\beta_1 \log^2 n) = O(d \log n)$. Since each column of Ψ is β-sparse, each column of Φ is s-sparse where $s = O(\beta_1 \log n)$. Finally, we show the algorithm $\mathcal{A}_{construct}$ for listing a non-zero entries in a column of Φ. Since, for each node $u \in L$, edges incident to u correspond to the non-zero entries of the u-th column of Ψ, and each edge can be computed in $\beta_2 = O(\text{poly}\log(n)))$ time, we have that all non-zero entries of each column of Φ can also be identified in $s\beta_2$ time. □

B.2.3 Proof of Theorem 4.16

Proof. First, we need this definition. An *Euler tour* $ET(T)$ of T is an (ordered) list of size $2|V(T)|-1$ with the following properties. Each element in the list is mapped to some node of T such that 1) each sublist of $ET(T)$ corresponds to some connected subtree of T, and 2) the whole list corresponds to the whole T. (See [16] for the precise definition of Euler tour.) For each node $v \in T$, there is some element from $ET(T)$ mapped to and we call the first element in $ET(T)$ mapped to v the *representative* of v in $ET(T)$. Given a tree T in F, an ET tree $\mathcal{T}_T$ is a balanced binary search tree whose leaves are exactly the elements of $ET(T)$. So each internal node u of $\mathcal{T}_T$, called *ET node*, corresponds to all leaves in the subtree rooted at u in $\mathcal{T}_T$, which form a sublist of $ET(T)$, which in turn corresponds to a subtree of T. For each $v \in T$, the leaf of $\mathcal{T}_T$ corresponding to the representative of v in $ET(T)$ is called a representative of v in $\mathcal{T}_T$.

We augment a standard ET-tree as follows. Each ET node u of $\mathcal{T}_T$ has an additional corresponding vector e_u of dimension d. Entries of e_u are either marked "active" or "inactive". We have the following invariant about e_u for each ET-node $u \in \mathcal{T}_T$: (1) if u is a leaf and u is not the representative of any vertex of T, then no entry of e_u is marked as active, (2) if u is a leaf and u is the representative of a vertex v of F, then, for $i \in [d]$, $e_u[i]$ is active if $x_v[i] \neq 0$ initially or $x_v[i]$ has been updated after preprocessing, and we have $e_u[i] = x_v[i]$ for each active entry, and (3) if u is an internal node with v and w as children in $\mathcal{T}_T$, then, for $i \in [d]$, $e_u[i]$ is active if either $e_v[i]$ or $e_w[i]$ is active, and the value of active entry $e_v[i]$ is the sum of $e_v[i]$ and $e_w[i]$ that are active.

It is easy to see that, given these invariants, the root r of $\mathcal{T}_T$ is such that if $e_r[i]$ is active, then $\sum_{u\in T} x_u[i] = e_r[i]$ and if $e_r[i]$ is inactive then $\sum_{u\in T} x_u[i] = 0$. Hence, to return return sum_vec, $\sum_{u\in T} x_u$ can be read off from e_r in time $O(d)$. Since ET-trees are balanced, the time for, update_vec, updating each some entry $x_u[i]$ is $O(\log n)$.

We analyze the time needed for preprocessing and other update operations. The preprocessing algorithm is to, for each T in F, mark as active and set the value of the entries of each ET node in $\mathcal{T}_T$ correctly according to the invariants. There are $nnz(x)$ entry to update, hence the preprocessing time takes $O(nnz(x)\log n)$. To update tree edges, by using the same algorithms as in typical ET

trees, each tree-edge update operation requires us to recompute the vector e_u of $O(\log n)$ many ET nodes u. So the time needed for insert_tree and delete_tree is $O(d \log n)$ time. □

B.3 From Section 4.3

B.3.1 Proofs of Lemma 4.23, Lemma 4.24, Lemma 4.25 and Lemma 4.26.

For completeness, we give the proofs of Lemma 4.23, Lemma 4.24, Lemma 4.25 and Lemma 4.26. In [16, Section 3.1], Henzinger and King show a reduction from dynamic k-weight MSF to dynamic SF problem. We modify the idea in their reduction and use it in all the proofs of this section. The basic idea is illustrated in the proof of Lemma 4.23 and Remark B.1.

Lemma B.8 (Reminder of Lemma 4.23)**.** *Suppose there is an (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph with preprocessing time t_p and update time t_u. Then there is a stable $(n, O(Tp_c), p_t, T)$-algorithm $\mathcal{B}$ for the same graph with preprocessing time $O(t_p + m_0)$ where m_0 is the number of initial edges, and update time $O(t_u \log n + \log^2 n)$.*

Proof of Lemma 4.23. The algorithm $\mathcal{B}$ preprocesses as follows. Given an initial graph G_0, preprocess G_0 using $\mathcal{A}$. Then $\mathcal{A}$ outputs a spanning forest F'. Then we find another spanning forest F of G_0 in time $O(m_0)$ where $m_0 = |E(G_0)|$. Then we initialize two augmented ET trees $\mathcal{E}$ and $\mathcal{E}'$. Both the input graph and forest of $\mathcal{E}'$ are F'. Similarly, both the input graph and forest of $\mathcal{E}$ are F. The total preprocessing time is $O(t_p + m_0)$.

To an update, we first describe how F' handle updates: We feed each update to $\mathcal{A}$ so that it updates its spanning forest F' of G in time t_u. As there are at most t_u changes in the tree edges of F', we update $\mathcal{E}'$ accordingly in time $O(t_u \log n)$.

To maintain F, if a tree-edge of F is not deleted, then F does not change. Suppose a tree-edge (u, v) is deleted and some connected component T of F is separated into $T_u \ni u$ and $T_v \ni v$. Then we query $\mathcal{E}'$ on F' using findtree operation, whether u and v are still connected in F'. If not, there there is no replacement edge to reconnect T_u and T_v and we are done. Otherwise, there is a path $P_{uv} = (u, \dots, v)$ in F' and there must be some edge (u', v') where $u' \in T_u$ and $v' \in T_v$. Using path operation on $\mathcal{E}'$ and findtree operation on $\mathcal{E}$, we can binary search in P_{uv} to find (u', v') and use it as a replacement edge in $O(\log n \times \log n) = O(\log^2 n)$ time. (To be more precise, we query $\mathcal{E}'$ for the k-th edge (u'', v'') in P_{uv} for some k, and check if $u'' \in T_u$ and $v'' \in T_v$ using $\mathcal{E}$). Therefore, the update time for F itself is $O(\log^2 n)$. Hence, the total update time of $\mathcal{B}$ is $O(t_u \log n + \log^2 n)$. By union bound, $\mathcal{B}$ is (n, p'_c, p'_t, U)-algorithm where $p'_c = O(Tp_c)$ and $p'_t = p_t$.

Finally, observe that $\mathcal{B}$ is stable because F is not updated except a tree-edge (u, v) of F is deleted. In that case, $\mathcal{B}$ either update a spanning forest into $F \setminus (u, v)$ or $F \cup (u', v') \setminus (u, v)$. □

Remark B.1. In the proof of Lemma 4.23, observe that get a stable algorithm $\mathcal{B}$ maintaining the spanning forest F, we just need to be able query connectivity between two nodes in some other spanning forest F' of G. Suppose that at any time i, we can query some spanning forest F'_i where F'_i can be totally different from F'_{i-1}. Then, by paying additional $O(m_0)$ in preprocessing time and $O(\log^2 n)$ in update time as in the proof of Lemma 4.23, we can also maintain F stably.

Next, we prove Lemma 4.24.

Lemma B.9 (Reminder of Lemma 4.24)**.** *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$ with preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then, for any $U \geq T$, there is a stable $(n, O(Up_c), O(p_t), U)$-algorithm $\mathcal{B}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$ with preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) + t_p(n, p_c, p_t, T)/T + \log^2 n)$.*

Proof of Lemma 4.24. Let G_i denote the graph after the i-th update. We divide the update sequence of length U for the algorithm $\mathcal{B}$ into phases, each of size $T_0 = T/2$. In high level, given the algorithm $\mathcal{A}$, we will first show an algorithm $\mathcal{A}'$ which uses two instances $\mathcal{A}_{odd}$ and $\mathcal{A}_{even}$ of $\mathcal{A}$ so that during each odd (even) phase, some spanning tree is maintained by $\mathcal{A}_{odd}$ ($\mathcal{A}_{even}$). However, the two spanning trees at the end of each odd phase and at the beginning of the next even phase can be totally different. After having $\mathcal{A}'$, we can devise the dynamic SF algorithm $\mathcal{B}$ as desired stated in Remark B.1. Let us denote $t_p = t_p(n, p_c, p_t, T)$ and $t_u = t_u(n, p_c, p_t, T)$ for convenience.

Now, we describe the algorithm $\mathcal{A}'$. In the first phase, we preprocess G using $\mathcal{A}_{odd}$ in time $t_p(n, p_c, p_t, T)$ and handle T_0 updates using $\mathcal{A}_{odd}$. For each non-first phase, we do the following. Consider each odd phase with time period $[i_0 + 1, i_0 + T_0]$. During time period $[i_0 + 1, i_0 + T_0/2]$, we preprocess G_{i_0} using $\mathcal{A}_{even}$ by evenly distributing the work so that each step takes at most $\frac{t_p}{T_0/2} = O(t_p/T)$ time. During time period $[i_0 + T_0/2, i_0 + T_0]$, at time $i_0 + T_0/2 + k$, we feed the two updates at time $i_0 + 2k - 1$ and $i + 2k$ to $\mathcal{A}_{even}$. Therefore, we have that after the $(i_0 + T_0)$-th update which is the beginning of even phase, $\mathcal{A}_{even}$ is maintaining a spanning tree of the current graph $G_{i_0+T_0}$. As $\mathcal{A}_{even}$ has only been only fed T_0 updates in time period $[i_0 + T_0/2, i_0 + T_0]$, $\mathcal{A}_{even}$ can handle $T - T_0 = T_0$ more incoming updates in the next even phase, taking time t_u on each update. We do symmetrically at each even phase. Therefore, we have that during each odd (even) phase, some spanning tree is maintained by $\mathcal{A}_{odd}$ ($\mathcal{A}_{even}$) as desired. The preprocessing time and update time of $\mathcal{A}'$ is t_p and $O(t_u + t_p/T)$ respectively.

After the i-th update, let F'_i be the spanning forest maintained by either $\mathcal{A}_{odd}$ or $\mathcal{A}_{even}$. We will assume that that an augmented ET tree from Theorem 3.3 is implemented on it. This slightly increase preprocessing time and update time of $\mathcal{A}'$ to $t_p + O(n) = O(t_p)$ and $O(t_u + t_p/T + n/T + \log n) = O(t_u + t_p/T + \log n)$. By Remark B.1, we can stably maintain a spanning forest F with the preprocessing time and update time $O(t_p + n) = O(t_p)$ and $O(t_u + t_p/T + \log^2 n)$ respectively.

By union bound, $\mathcal{B}$ is (n, p'_c, p'_t, U)-algorithm where $p'_c = O(Up_c)$ and $p'_t = O(p_t)$. To be more clear, we have $\mathcal{A}$ works correctly on all first U steps with probability $1 - Up_c$. Hence, $\mathcal{B}$ also works correctly on *all* first U steps with with probability $1 - Up_c$. In particular, $\mathcal{B}$ works correctly on each of the first U steps. For the update time, as we run two instances of $\mathcal{A}$ and other parts are deterministic, we have that $\mathcal{B}$ has preprocessing time and updated time as stated with probability $1 - 2p_t$. □

Lemma B.10 (Reminder of Lemma 4.25)**.** *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$, and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then there is a stable $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{B}$ for any graph with where the number of edges is always between $n/4$ and $2n$ with preprocessing time $O(t_p(O(n), p_c, p_t, T))$ and update time $O(t_u(O(n), p_c, p_t, T) + \log^2 n)$.*

Proof of Lemma 4.25. Let us define an (n, p_c, p_t, T)-algorithm $\mathcal{A}'$ for 2-weight MSF as an (n, p_c, p_t, T)-algorithm that maintains a 2-weight MSF instead of just a spanning forest. The reduction from dynamic SF to dynamic 2-weight MSF in [16, Section 3.1] can be stated as follows: given that there is $\mathcal{A}$, there is a stable $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{A}'$ for 2-weight MSF for a 3-bounded-degree graph where the number of edges is always between $n/4$ and $2n$, and $\mathcal{A}'$ has preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) + \log^2 n)$. Now, we show how to get $\mathcal{B}$ from $\mathcal{A}'$.

Let $G = (V, E)$ be an input graph for $\mathcal{B}$ where $n/2 \le |E| \le 2n$ but max degree of G maybe more than 3. Let G' be obtained from G by "splitting" each node in G into a path in G' of length equals to its degree. To be more precise, for each node $u \in V$ with incident edges $(u, v_1), \ldots, (u, v_s)$, we create a path $P_u = (u_1, \ldots, u_s)$ in G' and there are edges (u_i, v_i) for every i. The number of nodes in G' is still $O(n)$ because there are $O(n)$ edges in G. For each edge update of G, we can update G' accordingly using 3 edge updates.

Now, we assign weight 0 to edges in those paths, and 1 to original edges of G'. Observe that tree-edges with weight 1 of a spanning forest in G' form a spanning forest of G. Therefore, the algorithm $\mathcal{B}$ just runs a $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{A}'$ for 2-weight MSF on G' with preprocessing time $O(t_p(O(n), p_c, p_t, T))$ and update time $O(t_u(O(n), p_c, p_t, T) + \log^2 n)$ and maintain 2-weight minimum spanning forest F' in G'. By reporting only the updates of the tree-edges with weight 1 in the 2-weight MSF F' of G', we obtain updates for a spanning forest of G. □

Lemma B.11 (Reminder of Lemma 4.26)**.** *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph where the number of edges is always between $n/4$ and $2n$, and $\mathcal{A}$ has preprocessing time $t_p(n, p_c, p_t, T)$ and update time $t_u(n, p_c, p_t, T)$. Then there is a stable $(n, O(Tp_c), O(p_t), T)$-algorithm $\mathcal{B}$ for any graph with n nodes where the number of edges is always at most $2n$, and $\mathcal{B}$ has preprocessing time $O(t_p(O(m_0), p_c, p_t))$ where m_0 is a number of initial edges, and update time $O(t_u(O(m), p_c, p_t) + \log^2 n)$ where m is a number of edges when update.*

Proof. First, we define the notion of *reduced graphs*. Let $G_i = (V, E_i)$ be a graph after the i-th update for any i. Let V_i be the set of all endpoints of edges in E_i. The reduced graph G' of G_i is defined as follows: $G' = (V_i \cup X, E_i)$ where X is a set of $|E_i|$ *auxiliary nodes* which are isolated. So $|E_i| \leq |V(G')| \leq 2|E_i|$. After time i, suppose an edge (u, v) is inserted in G. If $u \notin V_i$, i.e. it is not an endpoint of edges in E_i, we let u' be some isolated node in X, add pointers between u and u', and say that $u' \in X$ is used. Otherwise $u \in V_i$, we let $u' = u$. We do similarly for v. Then we insert (u', v') into G'. When (u, v) is deleted in G, we delete (u', v') in G'. Suppose not all auxiliary nodes in X are used, and suppose that we can maintain a spanning forest in G', it is obvious how to maintain a spanning forest in G because there are pointers that maps between used node in X and nodes in G. Note also that we can easily construct the reduced graph G' of G_i in time $O(|E_i|)$ (as we assume the standard assumption that the list of edges of E_i is given).

We will use exactly the same technique as in the proof of Lemma 4.24. The update sequence is again divided into phases. We will use two instances $\mathcal{A}_{odd}$ and $\mathcal{A}_{even}$ of $\mathcal{A}$ so that during each odd (even) phase, some spanning tree is maintained by $\mathcal{A}_{odd}$ ($\mathcal{A}_{even}$). After having this, we can devise the stable dynamic SF algorithm $\mathcal{B}$ using augmented ET trees as in the proof of Lemma 4.24. However, in this reduction, the length of each phase is not fixed as in the proof of Lemma 4.24. Here, the number of edges in the current graph will define the length of the next phase. To be more precise, suppose that the current phase is a period $[i_0 + 1, i_0 + P_0]$. Then the length of the next phase is $|E(G_{i_0+P_0/2})|/4$. The length of the first phase is set to be $|E(G_0)|/2$. By this, we have the following invariant.

Claim B.12. For any phase of a period $[i_1 + 1, i_1 + P_1]$, we have $|E(G_{i_1})|/8 \leq P_1 \leq |E(G_{i_1})|/2$.

Proof. We prove by induction. Let $[i_0 + 1, i_0 + P_0]$ be the period of the previous phase, i.e. $i_1 = i_0 + P_0$. We claim that $|E(G_{i_0})| \leq 2|E(G_{i_1})|$. Otherwise $|E(G_{i_0})| - P_0 \leq |E(G_{i_0+P_0})| = |E(G_{i_1})| < |E(G_{i_0})|/2$, which means $P_0 > |E(G_{i_0})|/2$, contradicting the induction hypothesis.

By definition, we have $P_1 = \frac{|E(G_{i_0+P_0/2})|}{4} = \frac{|E(G_{i_1-P_0/2})|}{4}$. So we have

$$P_1 = \frac{|E(G_{i_1-P_0/2})|}{4} \leq \frac{|E(G_{i_1})| + P_0/2}{4} \leq \frac{|E(G_{i_1})| + |E(G_{i_0})|/4}{4} \leq \frac{|E(G_{i_1})| + |E(G_{i_1})|/2}{4} < |E(G_{i_i})|/2$$

,and we also have

$$P_1 = \frac{|E(G_{i_1-P_0/2})|}{4} \geq \frac{|E(G_{i_1})| - P_0/2}{4} \geq \frac{|E(G_{i_1})| - |E(G_{i_0})|/4}{4} \geq \frac{|E(G_{i_1})| - |E(G_{i_1})|/2}{4} = |E(G_{i_1})|/8$$

□

Now, we are ready to describe the algorithm. To preprocess the initial graph G_0, we construct a reduced graph G'_0 of G_0 and preprocess the reduced graph G'_0 using $\mathcal{A}_{odd}$ using $O(t_p(O(m_0), p_c, p_t) + m_0) = O(t_p(O(m_0), p_c, p_t))$ time in total where m_0 is a number of initial edges. Then we maintain a spanning forest in the reduced graph G'_0 using $\mathcal{A}_{odd}$ in the first phase.

For each odd phase with period $[i_0 + 1, i_0 + P_0]$, we maintain a spanning forest F_{odd} in the reduced graph using $\mathcal{A}_{odd}$. Note that, by Claim B.12, the number m' of edges in the reduced graph maintained by $\mathcal{A}_{odd}$ is always between $n'/4 \leq m' \leq n'$ where n' is the number of nodes in the reduced graph. This satisfies the condition of the underlying graph of $\mathcal{A}_{odd}$. Additionally, during time $[i_0 + \frac{P_0}{2} + 1, i_0 + \frac{3}{4}P_0]$, we evenly distribute the work for constructing a reduced graph $G'_{i_0+P_0/2}$ of $G_{i_0+P_0/2}$, and preprocessing $G'_{i_0+P_0/2}$ using $\mathcal{A}_{even}$. During time $[i_0 + \frac{3}{4}P_0 + 1, i_0 + P_0]$, we evenly feed the updates of period $[i_0 + \frac{P_0}{2} + 1, i_0 + P_0]$ to $\mathcal{A}_{even}$, so that at the beginning of even phase at time $i_0 + P_0 + 1$, $\mathcal{A}_{even}$ is maintaining a spanning tree of the current graph. As the number of edges during this phase is $\Theta(P_0)$ by Claim B.12, the update time in this odd phase is $O(t_p(O(P_0), p_c, p_t)/P_0 + t_u(O(P_0), p_c, p_t)) = O(t_p(O(m), p_c, p_t)/m + t_u(O(m), p_c, p_t)) = O(t_u(O(m), p_c, p_t))$ where m is the number of edges at the time of update. Note that $t_p(O(m), p_c, p_t) = O(m \cdot t_u(O(m), p_c, p_t))$ because we can always preprocess the graph of m edges by just updating each edge one-by-one. We work symmetrically in even phase. Therefore, we during each odd (even) phase, some spanning tree is maintained by $\mathcal{A}_{odd}$ ($\mathcal{A}_{even}$) as desired. By Remark B.1, we are done. The probability of failure is again bounded by union bound. □

B.3.2 Proofs of Lemma 4.27 and Lemma 4.28

In this section, we give the proofs of Lemma 4.27 and Lemma 4.28.

Lemma B.13 (Reminder of Lemma 4.27)**.** *Suppose there is a stable (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for a graph with n nodes where the number of edges is always at most $2n$, and $\mathcal{A}$ has preprocessing time $t_p(m_0, p_c, p_t, T)$ where m_0 is a number of initial edges and update time $t_u(m, p_c, p_t, T)$ where m is a number of edges when update. Suppose that for any n_1, n_2, $t_p(m_1, p_c, p_t, T) + t_p(m_2, p_c, p_t, T) \leq t_p(m_1 + m_2, p_c, p_t, T)$, then there is a stable $(n, O(np_c), O(np_t), T)$-algorithm $\mathcal{B}$ for any graph with n nodes with preprocessing time $O(t_p(m_0, p_c, p_t, T) \log n)$ where m_0 is a number of initial edges and update time $O(t_u(O(n), p_c, p_t, T))$*

Proof. We assume first that $\mathcal{A}$ is deterministic and we will remove the assumption later. We partition a complete graph with n nodes into n arbitrary trees $T_1, \ldots, T_n$. For example, T_i can be a star having the vertex i as a core and vertices $1, \ldots, i-1$ as leaves. Then we define a complete binary tree called *sparsification tree* $\mathcal{T}$ with n leaves corresponding to $T_1, \ldots, T_n$. For each node $x \in \mathcal{T}$, let $\mathcal{T}_x$ be the subtree of $\mathcal{T}$ rooted at x, and x corresponds to a graph H_x formed by all edges of T_i where T_i are the leaves in $\mathcal{T}_x$. The root of $\mathcal{T}$ corresponds to the complete graph.

Let G be the dynamic graph which is an input of the algorithm $\mathcal{B}$. For each node $x \in \mathcal{T}$, we define G_x to be a graph such that $E(G_x) = E(G) \cap E(H_x)$. For every node x, let G'_x be a certain subgraph (to be define next) of G_x which is an underlying graph of an instance of the algorithm $\mathcal{A}$. Let F_x be the spanning forest of G'_x maintained by $\mathcal{A}$. The following invariant about G'_x must hold: if x is a leaf of $\mathcal{T}$, then we set $G'_x = G_x$. If v is an internal node of $\mathcal{T}$ with y and z as children, then $G'_x = F_y \cup F_z$. Observe that for any $x \in \mathcal{T}$, $|E(G'_x)| \leq 2n$. Observe further that F_x is not only a spanning forest of G'_x but, by induction, it is also a spanning forest of G_x. Hence, at the root r of $\mathcal{T}$, F_r is a spanning forest of G as desired.

Now, we are ready to describe how $\mathcal{B}$ works. $\mathcal{B}$ preprocesses the initial graph G as follows. First, for each $x \in \mathcal{T}$, since G'_x is an underlying graph of an instance of the algorithm $\mathcal{A}$, $\mathcal{A}$ requires that G'_x

is represented as a list of its edge. We first do this by scanning edges in G, and obtain the list of edges on G'_x for all leaves $x \in \mathcal{T}$. Then we preprocess those G'_x using the algorithm $\mathcal{A}$ and obtain F_x. For each non-leaf node $x \in \mathcal{T}$, we proceed in bottom-up manner: let $y, z \in \mathcal{T}$ be the children of x, we scan the list of edges in $F_y \cup F_z = G'_x$ and preprocess G'_x and obtain F_x. To bound the running time, let $\mathcal{T}^d$ be a set of nodes in $\mathcal{T}$ with depth d. For different nodes $x, y \in \mathcal{T}^d$, we know that G'_x and G'_y are edge disjoint. So $\sum_{x\in\mathcal{T}^d} |E(G'_x)| \leq m$, and hence the preprocessing time spent on graphs G'_x, for all $x \in \mathcal{T}^d$, is $\sum_{x\in\mathcal{T}^d}(t_p(n, |E(G'_x)|, p_c, p_t, T)) + O(|E(G'_x)|) = O(t_p(n, m_0, p_c, p_t, T))$ by the assumption about $t_p(\cdot)$. Since $\mathcal{T}$ has depth $O(\log n)$, the total preprocessing time is $O(t_p(n, m_0, p_c, p_t, T) \log n)$.

Next, we describe how to update. The goal is to maintain the invariant that $G'_x = G_x$ for each leaf $x \in \mathcal{T}$ and $G'_x = F_y \cup F_z$ for each non-leaf $x \in \mathcal{T}$ with children y and z. When an edge e is updated in G, let x be the unique leaf of $\mathcal{T}$ that H_x contains e. Let $(x = x_0, x_1, \ldots, x_s = r)$ be the path in $\mathcal{T}$ from x to the root r of $\mathcal{T}$. There are two cases.

First, suppose the update is to insert $e = (u, v)$. Let k be the smallest index such that u and v were connected in G'_{x_k} before the insertion. For $j = 1$ to $k-1$, u and v were *not* connected in G'_{x_j}, so e will be added into F_{x_j} and this is the only change in F_{x_j} as F_{x_j} is stable. So only e is inserted into $G'_{x_{j+1}}$ in the next level. For $j = k$, u and v were connected in G'_{x_j} and e will be inserted and become a non-tree edge in G'_{x_j} as F_{x_j} is stable. So there is no further update to the next level. Therefore, we only do at most one update on $O(\log n)$ graphs, which takes $O(t_u(n, 2n, p_c, p_t, T) \log n)$ time in total.

Second, suppose the update is to delete $e = (u, v)$. Let k be the largest index such that u and v were a tree-edge of F_{x_k} before the deletion. So e is also a non-tree edge in $G'_{x_{k+1}}$ and e is not in any G'_{x_j} for $j > k + 1$. First, we delete the edge e from all $G'_{x_0}, \ldots, G'_{x_{k+1}}$ using $\mathcal{A}$ and find the replacement edge in each graph if exists. Let ℓ be the smallest index such that u and v are still connected in G'_{x_ℓ}after the deletion. For $j < \ell$, there is no replacement edge. So removing e from F_{x_j} is the only change by stability of F_{x_j}. To maintain the invariant, there is no further update needed to be propagate to $G'_{x_{j+1}}$, since e is already removed from $G'_{x_{j+1}}$. For j such that $\ell \leq j \leq k$, there is also a replacement edge e_j newly added into F_{x_j}. By stability of F_{x_j}, only e is removed and e_j is added into F_{x_j}. Since e is already removed from $G'_{x_{j+1}}$, we only need to insert e_j into $G'_{x_{j+1}}$. But we know that $e_j \notin F_{x_{j+1}}$ because u and v are already reconnected in $F_{x_{j+1}}$ by e_{j+1}. Hence, the insertion of e_j into $G'_{x_{j+1}}$ will not be propagated into the next level $j + 2$. For $j = k + 1$, we know that e is not a tree-edge. So $F_{x_{k+1}}$ does not change and also and so there is no further update in the upper level. Therefore, we only do at most two update on $O(\log n)$ graphs, which takes $O(t_u(n, 2n, p_c, p_t, T) \log n)$ time in total.

Finally, to remove the assumption that $\mathcal{A}$ is deterministic, we know that if all instances of $\mathcal{A}$ works correctly on G'_x for all $x \in \mathcal{T}$, then F_r is a spanning forest of G where r is the root of $\mathcal{T}$. There are $O(n)$ nodes in $\mathcal{T}$ and each instance works correctly at any time with probability $1 - p_c$. By union bound F_r is a spanning forest of G with probability $1 - O(np_c)$. By the same argument, as there are n instances of $\mathcal{A}$, the preprocessing time and update time of $\mathcal{B}$ are bounded as state with probability $1 - O(np_t)$.

We note that here, we only show that the algorithm $\mathcal{B}$ has update time $O(t_u(n, 2n, p_c, p_t, T) \log n)$. However, a logarithmic factor can be shaved off using a more clever sparsification tree $\mathcal{T}$ as described in [11, Section 3.1] where each node has 3 or 4 leaves and the number of vertices at level $2i$ is at most $O(n/2^i)$. □

Lemma B.14 (Reminder of Lemma 4.28)**.** *Suppose there is an (n, p_c, p_t, T)-algorithm $\mathcal{A}$ for any graph with n nodes with preprocessing time $t_p(m_0, p, p_t, T)$ where m_0 is a number of initial edges and update time $t_u(n, p_c, p_t, T)$. If $T \geq n^2$, then there is an $(n, O(Tp_c), O(p_t), \infty)$-algorithm $\mathcal{B}$ for any graph with n nodes with preprocessing time $O(t_p(m_0, p_c, p_t, T))$ where m_0 is a number of initial*

edges and update time $O(t_u(n, p_c, p_t, T))$.

Proof. The reduction from dynamic SF to dynamic 2-weight MSF in [16, Section 3.1] can be stated as follows: given that there is $\mathcal{A}$, there is a stable (n, p'_c, p'_t, T)-algorithm $\mathcal{A}'$ for 2-weight MSF for any graph where $p'_c = O(Tp_c)$ and $p'_t = O(p_t)$, and $\mathcal{A}'$ has preprocessing time $O(t_p(n, p_c, p_t, T))$ and update time $O(t_u(n, p_c, p_t, T) + \log^2 n)$. We will use $\mathcal{A}'$ in the reduction below. Let $\mathcal{A}^{odd}$ and $\mathcal{A}^{even}$ be two instances of the algorithm $\mathcal{A}$. We denote F^{odd}, F^{even} and F' as the spanning forests maintained by $\mathcal{A}^{odd}$, $\mathcal{A}^{even}$ and $\mathcal{A}'$ respectively. For any time i, F_i^{odd}, F_i^{even} and F'_i are F^{odd}, F^{even} and F' after the i-th update respectively. We divide the sequence of updates in phases. Each phase has size $P = T/2$. We will first show the reduction assuming that $\mathcal{A}^{odd}$, $\mathcal{A}^{even}$ and $\mathcal{A}'$ are deterministic, and then we will remove the assumption later. For convenience, we denote $t_p(m) = t_p(m_0, p_c, p_t, T)$ for any m.

Suppose that i_0 is the ending time of some odd phase. For each update at time $i_0 + 1$ to $i_1 = i_0 + 2t_p(n^2)$, we feed each update to $\mathcal{A}^{odd}$ taking time $t_u(n)$ per update and then report the answer from $\mathcal{A}^{odd}$ (the list of tree-edges to be added or removed from the spanning forest F^{odd} maintained by $\mathcal{A}^{odd}$). But, in addition, during this period of time we will also spend time on $\mathcal{A}^{even}$ as follows. We distribute evenly, from time i_0+1 to $i_0+t_p(n^2)$, the work of $\mathcal{A}^{even}$ for preprocessing the graph G_{i_0} at time i_0 which has at most n^2 edges. This takes $O(1)$ time per update. So, at time $i_0 + t_p(n^2)$, $\mathcal{A}^{even}$ finishes preprocessing G_{i_0} which is a graph from $t_p(n^2)$ steps before. For each time $i_0 + t_p(n^2) + k$ where $k \in [1, t_p(n^2)]$, we feed the two updates from time $i_0 + 2k - 1$ and i_0+2k. This takes $O(t_u(n))$ time per update. Hence, at time $i_1 = i_0+2t_p(n^2)$, $\mathcal{A}^{even}$ is maintaining a spanning forest $F_{i_1}^{even}$ of the current graph G_{i_1}.

Next, for each update at time $i_1 + 1$ to $i_2 = i_1 + t_p(0) + n$, we feed each update to both $\mathcal{A}^{odd}$ and $\mathcal{A}^{even}$ taking time $O(t_u(n))$ per update, but we only report the answer from $\mathcal{A}^{odd}$. During this period of time we will, in addition, spend time on $\mathcal{A}'$. From time $i_1 + 1$ to $i_1 + t_p(0)$, we distribute evenly the work of $\mathcal{A}'$ for preprocessing an empty graph.[18] At time $i_1 + t_p(0)$, $\mathcal{A}'$ is ready to handle an update. Let G' be the underlying graph of $\mathcal{A}'$. Now, the goal is that at time i_2 we have edges of G' is exactly F^{odd} (i.e. $G'_{i_2} = F_{i_2}^{odd}$), and each edge has weight 1. To do this, for each time $i_1 + t_p(0) + 1$ to $i_2 = i_1 + t_p(0) + n$, we have the invariant that $E(G') \subseteq F^{odd}$. We insert a constant number, say three, of edges of weight 1 from F^{odd} to G' and delete an edge e from G' whenever e is removed from F^{odd}. Since F^{odd} is stable, at each step there is at most one edge removed from F^{odd}. But because we insert strictly more than one edge at each step, and F^{odd} has at most $n - 1$ edges, and we spend at most n steps doing this, taking time $O(t_u(n))$ per step. We conclude $G'_{i_2} = F_{i_2}^{odd}$. Note that we have $F'_{i_2} = F_{i_2}^{odd}$ as well.

Next, from each update at time $i_2 + 1$ to $i_3 = i_2 + n$, we feed each update to all $\mathcal{A}^{odd}$, $\mathcal{A}^{even}$ and $\mathcal{A}'$. But we use report the answer from $\mathcal{A}'$. Note that we can switch from reporting the answer from $\mathcal{A}^{odd}$ to reporting from $\mathcal{A}'$ because $F'_{i_2} = F_{i_2}^{odd}$. Now, the goal is that at time i_3 we have $G'_{i_3} = F_{i_3}^{odd} \cup F_{i_3}^{even}$ and each edge $e \in F^{even}$ has weight 0 in G'. Therefore, we have $F'_{i_3} = F_{i_3}^{even}$. Even when having the invariant that $F^{odd} \subseteq E(G') \subseteq F^{odd} \cup F^{even}$, this can still be done similarly as above by doing a constant number of updates to G' at each step, taking time $O(t_u(n))$ per update. At time i_3, we now have $F'_{i_3} = F_{i_3}^{even}$ and we terminate $\mathcal{A}^{odd}$ and $\mathcal{A}'$ (they will be initialized again around the ending time of the next phase).

Next, from time i_3+1 to i_0+P which is the ending time of next phase, we just feed each update to $\mathcal{A}^{even}$ and also report its answer. Note again that we can switch from reporting the answer from $\mathcal{A}'$ to reporting from $\mathcal{A}^{even}$ because $F'_{i_3} = F_{i_2}^{even}$.

We need that $i_3 + 1 \leq i_0 + P$. This is true because $i_3 - i_0 + 1 \leq 2t_p(n^2) + t_p(0) + n + n + 1 \leq$

[18] Note that $t_p(0)$ may not be trivial. For instance, $t_p(0)$ can be linear in n.

$6t_p(n^2) \leq P$. We also need that $\mathcal{A}^{odd}$ and $\mathcal{A}'$ need to handle at most T updates before they are terminated. This is true because $\mathcal{A}^{odd}$ is surely initialized after time $i_0 - P$ and terminated before time $i_0 + P$, and $T = 2P$. Similar argument holds for $\mathcal{A}^{even}$ and $\mathcal{A}'$.

To conclude, we have obtained an algorithm $\mathcal{B}$ that alternatively invoke $\mathcal{A}^{odd}$ and $\mathcal{A}^{even}$ on each phase, and periodically invoke $\mathcal{A}'$ during the transition between each phase. The preprocessing time is $t_p(m)$, which is exactly preprocessing time of $\mathcal{A}^{odd}$. The update time is $O(t_u(n))$ time.

Now, we remove the assumption that $\mathcal{A}^{odd}$, $\mathcal{A}^{even}$ and $\mathcal{A}'$ are deterministic, and we want to prove that the probability at the maintained forest by $\mathcal{B}$ spans the underlying graph at each step with probability at least $1 - O(Tp_c)$. Observe that, at any time, the forest $F^{\mathcal{B}}$ maintained by B is the forest maintained by either $\mathcal{A}^{odd}$, $\mathcal{A}^{even}$ or $\mathcal{A}'$, i.e. F^{odd}, F^{even}, or F' respectively. From the algorithm, when we report either F^{odd} or F^{even}, the underlying graph is G itself. So F^{odd} and F^{even} are both spanning forests of G with probability at least $1 - p_c$. However, when we report F', the underlying graph is G' where $F^{odd} \subseteq E(G') \subseteq F^{odd} \cup F^{even}$. Since F^{odd} and F^{even} are spanning forests of G with probability $1-p_c$ and F' is a spanning forest of G' with probability $1-p'_c$, we have F' is a spanning forest of G with probability at least $1 - p'_c - 2p_c = 1 - O(Tp_c)$ by union bound. Using union bound, we have that the failure probability about running time is $O(p_t)$. □

B.4 Proof of Theorem 5.4

We prove Theorem 5.4 by extending the cut-matching game of [26]. In the following, the notion of graph embedding is required. We say that a weighted graph $H = (V, E_H, w_H)$ can be *embedded* in another weighted graph $G = (V, E_G, w_G)$ with congestion C iff a flow of $w_H(f)$ units can be routed in G between the end-points of f *simultaneously* for all $f \in E_H$ without violating the edge-capacities $w_G(e)$, for any edge $e \in E_G$, by a factor more than C. we state an easy observation:

Proposition B.15. *Let G and H be two unweighted undirected graphs with the same set of vertices. If H can be embedded in G with congestion C, then $\phi(G) \geq \phi(H)/C$. Moreover, if every cut of H of size at least s has expansion at least α, then every cut of G of size at least s has expansion at least α/C.*

Next, we prove some lemmas which are extending the cut-matching game of [26].

Lemma B.16. *Given an undirected graph $G = (V, E)$ with n vertices and m edges, an expansion parameter $\alpha > 0$, and an approximation ratio $\gamma \geq 1$, there is a randomized algorithm that with high probability outputs either*

- *an α-sparse cut S where $|S| \leq n/2$ or*
- *a graph H embeddable in G with congestion at most $O(\gamma \log^3 n/\alpha)$ where $\phi(H) \geq 1/2$.*

Furthermore the algorithm runs in $O(\gamma \log^4 n \cdot T(m))$ time where $T(m)$ is the time for computing γ-approximate max-flow in an undirected graph with m edges.

Proof. We first summarize the cut-matching game [26] that our proof is based on. Given an expansion parameter α on a graph $G = (V, E)$, there are two players, a cut player and a matching player, interacting with each other for $r = \Theta(\log^2 n)$ rounds. In the i-th round, the cut player will choose a cut $(S_i, V \setminus S_i)$ of size $|S_i| = n/2$, and then the matching player either 1) outputs a perfect matching M_i between S_i and $V \setminus S_i$ and the game is continued to the next round, or 2) outputs an α-sparse sparse cut and terminate the game.

It is shown in [26] that, after r rounds, if the matching player did not output an α-sparse cut, then a union of the matchings $H = \bigcup_{i=1}^{r} M_i$ is such that $\phi(H) \geq 1/2$. Moreover, the time needed for

the cut player in each round is only $O(n)$. For the matching player, by computing *exact* max flow in each round, the matching M_i found by the matching player can be embedded in G with congestion at most c/α where $c = 1$. Hence, H can be embedded in G with congestion $r \times c/\alpha = O(\log^2 n/\alpha)$. However, computing *exact* max flow is too costly for us.

We will show that if we relax $c = \Theta(\gamma \log n)$, then the matching player only needs to compute c many γ-approximate max flow computations to finish each round. On the high level, in the i-th round of matching player, given a cut $(S_i, V \setminus S_i)$ from the cut player, there will be c sub-rounds. In each sub-round, the matching player will compute a γ-approximate max flow. He will either output an α-sparse cut, or, after c sub-rounds, output a perfect matching to complete the i-th round for the cut-matching game. We describe the procedure precisely below.

In the j-th sub-round, let $T_1^{(j)}, T_2^{(j)} \subset V$ be two disjoint sets of vertices each of size $k^{(j)} = |T_1^{(j)}| = |T_2^{(j)}|$. In the first sub-round, $T_1^{(1)} = S_i$ and $T_2^{(1)} = V \setminus S_i$. Let $G^{(j)}$ be a *weighted* undirected graph obtained from G by adding two new vertices t_1 and t_2. There are edges in $G^{(j)}$, with unit capacity, from t_1 to every vertex in $T_1^{(j)}$, and similarly from t_2 to $T_2^{(j)}$. The original edges of G have capacity $1/\alpha$ in $G^{(j)}$. The matching player runs a γ-approximate max flow algorithm in $G^{(j)}$, which also produces an γ-approximate min cut $C^{(j)}$.

Claim B.17. If $\delta_{G^{(j)}}(C^{(j)}) < k^{(j)}$ then $\phi_G(C^{(j)} \setminus \{t_1\}) < \alpha$.

Proof. This is similar to [26, Lemma 3.7]. Let n_{t_1} and n_{t_2} be a number of cut edges $\partial_{G^{(j)}}(C^{(j)})$ incident to t_1 and t_2 respectively. Observe that $\delta_{G^{(j)}}(C^{(j)}) = \frac{1}{\alpha}\delta_G(C^{(j)} \setminus \{t_1\}) + n_{t_1} + n_{t_2}$. So $\delta_G(C^{(j)} \setminus \{t_1\}) = \alpha(\delta_{G^{(j)}}(C^{(j)}) - n_{t_1} - n_{t_2})$. But $|C^{(j)} \setminus \{t_1\}| \geq k^{(j)} - n_{t_1} > \delta_{G^{(j)}}(C^{(j)}) - n_{t_1}$ and similarly we can show that $|V \setminus \{C^{(j)} \cup \{t_2\}\}| > \delta_{G^{(j)}}(C^{(j)}) - n_{t_2}$, therefore, $\phi_G(C^{(j)} \setminus \{t_1\}) = \frac{\delta_G(C^{(j)} \setminus \{t_1\})}{\min\{|C^{(j)} \setminus \{t_1\}|, |V \setminus \{C^{(j)} \cup \{t_2\}\}|\}} < \alpha$. □

Hence, if the total weight of the approximate min cut $C^{(j)}$ is $\delta_{G^{(j)}}(C^{(j)}) < k^{(j)}$, then the matching player can return the smaller side of $(C^{(j)} \setminus \{t_1\}, V \setminus \{C^{(j)} \cup \{t_2\}\})$ which is the desired α-sparse cut, and terminate the cut-matching game. Otherwise, we know that the size of the max flow is at least $k^{(j)}/\gamma$, and by [21] the flow rounding algorithm in $O(m \log(n^2/m))$ time, we can assume that the flow is integral. We iteratively decompose the flow into paths using dynamic tree [36] in total time $O(m \log n)$. Then we construct a partial matching $M_i^{(j)}$ between $T_1^{(j)}$ and $T_2^{(j)}$ of size at least $k^{(j)}/\gamma$ as follows: For each decomposed path P, if $P = (t_1, u_1, \ldots, u_2, t_2)$ where $u_1 \in T_1^{(j)}$ and $u_2 \in T_2^{(j)}$ then we add $(u_1, u_2) \in M_i^{(j)}$. There are at least $k^{(j)}/\gamma$ paths because of the size of the max flow, hence the size of $M_i^{(j)}$ is at least $k^{(j)}/\gamma$. Let $T_1^{(j+1)} = T_1^{(j)} \setminus V(M_i^{(j)})$ and $T_2^{(j+1)} = T_2^{(j)} \setminus V(M_i^{(j)})$. As $|T_i^{(j+1)}| \leq (1 - \frac{1}{\gamma})|T_i^{(j)}|$, there are at most $c = \Theta(\gamma \log n)$ sub-rounds before $T_1^{(j+1)}, T_2^{(j+1)} = \emptyset$. If the game is not terminated until then, the matching player returns a perfect matching $M_i = \bigcup_j M_i^{(j)}$ between S_i and $V \setminus S_i$. Observe that M_i can be embedded in G with congestion c/α as desired, because each $M_i^{(j)}$ can be embedded in G with congestion $1/\alpha$ as certified by the paths from the flow decomposition.

To conclude, when we relax $c = O(\gamma \log n)$, we can simulate one round of matching player in the cut matching game of [26] using c sub-rounds where each sub-round needs just one γ-approximate max flow computation. As a result, we either get an α-sparse cut, or a graph $H = \bigcup_{i=1}^{r} M_i$ embeddable in G with congestion $O(r \times c/\alpha) = O(\gamma \log^3 n/\alpha)$ where $\phi(H) \geq 1/2$. The total cost for the cut player is $r \times O(n)$. The total cost for the matching player is $rc \times O(T(m) + m \log n)$. So the total running time is $O(\gamma \log^4 n \cdot T(m))$.

□

Next, we extend the previous proof so that we can compute a sparse cut of any specified size s.

Lemma B.18. *Given an undirected graph $G = (V, E)$ with n vertices and m edges, an expansion parameter $\alpha > 0$, an approximation ratio $\gamma \geq 1$, and a size parameter $s = O(\frac{n}{\log^2 n})$, there is a randomized algorithm that with high probability outputs either*

- *an α-sparse cut S where $s \leq |S| \leq n/2$, or*
- *a graph H embeddable in G with congestion at most $O(\gamma \log^3 n/\alpha)$ where every cut in H of size $\Omega(s \log^2 n)$ has expansion at least $1/4$.*

Furthermore the algorithm runs in $O(\gamma \log^4 n \cdot T(m))$ time where $T(m)$ is the time for computing γ-approximate max-flow in a graph with m edges.

Proof. We extend the proof of Lemma B.16. Suppose that we run the cut matching game in the i-th round and in the j-subround of matching player. We are going to find a partial matching between two disjoint set $T_1^{(j)}, T_2^{(j)} \subset V$ each of size $k^{(j)} = |T_1^{(j)}| = |T_2^{(j)}|$. Let $C^{(j)}$ be the γ-approximate min cut computed by a γ-approximate max flow algorithm in a graph $G^{(j)}$, which defined as in the proof of Lemma B.16.

Claim B.19. If $\delta_{G^{(j)}}(C^{(j)}) \leq k^{(j)} - s$, then two sides of the cut are of size at least s i.e. $|C^{(j)} \setminus \{t_1\}|, |V \setminus (C^{(j)} \cup \{t_2\})| \geq s$.

Proof. We only give a proof for $|C^{(j)} \setminus \{t_1\}|$ because they are symmetric. Let n_{t_1} and n_{t_2} be a number of cut edges $\partial_{G^{(j)}}(C^{(j)})$ incident to t_1 and t_2 respectively. As in Claim B.17, $\delta_{G^{(j)}}(C^{(j)}) = \frac{1}{\alpha}\delta_G(C^{(j)} \setminus \{t_1\}) + n_{t_1} + n_{t_2} \geq n_{t_1}$ and hence $|C^{(j)} \setminus \{t_1\}| \geq k^{(j)} - n_{t_1} \geq k^{(j)} - \delta_{G^{(j)}}(C^{(j)}) \geq s$. □

So, if $\delta_{G^{(j)}}(C^{(j)}) \leq k^{(j)} - s$, then we return the smaller side of $(C^{(j)} \setminus \{t_1\}, V \setminus \{C^{(j)} \cup \{t_2\}\})$ which has expansion less than α by Claim B.17 and has size at least s by Claim B.19. Otherwise, the size of the max flow is at least $(k^{(j)} - s)/\gamma$. Again, we can decompose the flow into paths which define a partial matching $M_i^{(j)}$ between $T_1^{(j)}$ and $T_2^{(j)}$ of size at least $(k^{(j)} - s)/\gamma$. Let $T_1^{(j+1)} = T_1^{(j)} \setminus M_i^{(j)}$ and $T_2^{(j+1)} = T_2^{(j)} \setminus M_i^{(j)}$. We know that $(|T_i^{(j+1)}| - s) \leq (1 - \frac{1}{\gamma})(|T_i^{(j)}| - s)$. So there are at most $c = \Theta(\gamma \log n)$ sub-rounds before $|T_1^{(j+1)}|, |T_2^{(j+1)}| \leq s$. If the game is not terminated until then, let M_i'' be an *arbitrary* matching between $T_1^{(j+1)}$ and $T_2^{(j+1)}$, and return a perfect matching $M_i' = \bigcup_j M_i^{(j)} \cup M_i''$ between S_i and $V \setminus S_i$ as an output of matching player in the i-th round.

Suppose that, after r rounds, the matching game did not return a sparse cut. Let $H' = \bigcup_{i=1}^r M_i'$ and $H'' = \bigcup_{i=1}^r M_i''$. We claim that $H = H' \setminus H''$ is our desired output. Observe that $H = \bigcup_{i=1}^r (\bigcup_j M_i^{(j)})$, so H can be embedded in G with congestion $O(rc/\alpha) = O(\gamma \log^3 n/\alpha)$. It is left to show that every cut in H of size $4rs = \Omega(s \log^2 n)$ has expansion at least $1/4$:

Claim B.20. Let $S \subset V$ be a cut of size $4rs \leq |S| \leq n/2$, then $\phi_H(S) \geq 1/4$.

Proof. From the property of cut matching game shown in [26], H' is such that $\phi(H') \geq 1/2$. So $\delta_{H'}(S) \geq |S|/2$. Since $\delta_H(S) = \delta_{H'}(S) - \delta_{H''}(S)$ and the number of edges in $H'' = \bigcup_{i=1}^r M_i''$ is at most rs, we conclude that $\delta_H(S) = \delta_{H'}(S) - \delta_{H''}(S) \geq |S|/2 - rs \geq |S|/2 - |S|/4 = |S|/4$. □

The running time analysis is as in Lemma B.16. □

By plugging in the $\tilde{O}(m)$-time $(1+\epsilon)$-approximate max flow algorithm by [33] to Lemma B.16 and Lemma B.18, and by assigning the parameter c_{size} and c_{exp}, we get the following corollaries:

Corollary B.21. *Given an undirected graph $G = (V, E)$ with n vertices and m edges, an expansion parameter $\alpha > 0$, there is a $\tilde{O}(m)$-time randomized algorithm that with high probability outputs either*

- *an α-sparse cut S where $|S| \leq n/2$ or*
- *a graph H embeddable in G with congestion at most c'_{exp}/α where $\phi(H) \geq 1/2$,*

where $c'_{exp} = \Theta(\log^3 n)$.

Corollary B.22. *Given an undirected graph $G = (V, E)$ with n vertices and m edges, an expansion parameter $\alpha > 0$, and a size parameter $s \leq n/c'_{size}$, there is a $\tilde{O}(m)$-time randomized algorithm that with high probability outputs either*

- *an α-sparse cut S where $s \leq |S| \leq n/2$, or*
- *a graph H embeddable in G with congestion at most c'_{exp}/α where every cut in H of size at least sc_{size} has expansion at least $1/4$.*

where $c'_{size} = \Theta(\log^2 n)$ and $c'_{exp} = \Theta(\log^3 n)$.

Now, we are ready to prove Theorem 5.4.

Proof of Theorem 5.4. Given a graph $G = (V, E)$ with parameter α, Let c'_{size} and c'_{exp} be the parameters from Corollary B.21 and Corollary B.22. Set $c_{size} = 2c'_{size}$ and $c_{exp} = 4c'_{exp}$ to be the parameter of Theorem 5.4. Let A and B be the algorithms from Corollary B.21 and Corollary B.22, respectively.

We give G and α to A. If A returns a graph H embeddable in G with congestion c'_{exp}/α where $\phi(H) \geq 1/2$. By Proposition B.15, we can just report $\phi(G) \geq \phi(H)2\alpha/c'_{exp} \geq \alpha/c_{exp}$ and we are done.

Suppose that A returns an α-sparse cut S_A where $|S_A| \leq n/2$. We binary search for the largest s such that when B is given G, α and s as inputs, then B returns an α-sparse cut. Suppose that, when $s = s^* - 1$, B returns an α-sparse cut S_B where $s^* - 1 \leq |S_B| \leq n/2$, and when $s = s^*$, B returns a graph H_B embeddable in G with congestion c'_{exp}/α where every cut in H_B of size at least $c'_{size}s^*$ has expansion at least $1/4$. By Proposition B.15, this implies that every cut S of size at least $c'_{size}s^*$ in G has expansion at least $\alpha/4c'_{exp} = \alpha/c_{exp}$. Therefore, $\mathsf{OPT}(G, \alpha/c_{exp}) < c'_{size}s^* = c_0 s^*/2$. We can return S_B because $|S_B| = s^* - 1 \geq s^*/2 > \mathsf{OPT}(G, \alpha/c_{exp})/c_{size}$. Otherwise, suppose that B never returns a cut.

Then, when $s = 1$, B returns a graph H_B embeddable in G with congestion c'_{exp}/α where every cut in H_B of size at least c'_{size} has expansion at least $1/4$. By Proposition B.15, this implies that every cut S of size at least c'_{size} in G has expansion at least $\alpha/4c'_{exp} = \alpha/c_{exp}$. Therefore, $\mathsf{OPT}(G, \alpha/c_{exp}) < c'_{size} \leq c_{size}$. So we can return S_A because $|S_A| \geq 1 > \mathsf{OPT}(G, \alpha/c_{exp})/c_{size}$. □

B.5 Proof of Theorem 6.5

First, we slightly change the definition of augmented graphs defined in [30].

Definition B.23 (Augmented Graph)**.** *Given an undirected graph $G = (V, E)$, a set $A \subset V$ where $|A| \leq |V - A|$, an expansion parameter α, and an overlapping parameter $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$, the* augmented graph $G_A(\alpha, \sigma)$ *is the capacitated directed graph with the following properties:*

- *The vertex set* $V(G_A(\alpha,\epsilon)) = V(G) \cup \{s,t\}$. *We call* s *and* t *the source and the sink vertices, respectively.*
- *The edge set* $E(G_A(\alpha,\epsilon))$ *contains all original undirected*[19] *edges* $e \in E(G)$ *with capacity* $1/\alpha$. $E(G_A(\alpha,\epsilon))$ *also contains directed edges* (s,u) *for all* $u \in A$ *with unit capacity and* (v,t) *for all* $v \in V - A$ *with capacity* ϵ_σ *where* $\epsilon_\sigma \triangleq \frac{1}{3(1/\sigma-1)} \in [\frac{|A|}{|V-A|}, 1]$.

The following lemma about the properties of s-t min cut in augmented graphs illustrates the usefulness of augmented graphs. The proof below is different from Lemma 3.2 in [30] in two aspects. First, we argue about the expansion of cuts instead of arguing about conductance of cuts as in [30], which is a slight different notion. Second, we show that the s-t min cut is a most-balanced sparse cut, not just a sparse cut as in [30].

Note that s-t max-flow-min-cut value of the augmented graph $G_A(\alpha,\sigma)$ is at most $|A|$ because the cut $\{s\}$ has capacity $|A|$.

Lemma B.24. *Given an undirected graph* $G = (V,E)$, *a set* $A \subset V$ *where* $|A| \le 4|V-A|$, *an expansion parameters* α *and an overlapping parameter* $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$, *let* F^* *be the* s-t *max-flow-min-cut value of the augmented graph* $G_A(\alpha,\sigma)$. *Suppose that* $F^* = |A| - c$ *for some value* $c \ge 0$. *Then we have that following:*

- *For any* (A,σ)*-overlapping cut* S *in* G *where* $|S| \ge 3c/\sigma$, *then* $\phi(S) \ge \frac{\sigma}{3}\alpha$, *and*
- *if* $c > 0$, *there is an* α*-sparse cut* S' *in* G *such that* $|S'|, |V - S'| \ge c$.

Proof. Let $G' = G_A(\alpha,\sigma)$. For any cut S, let $\delta'(S)$ be the capacity of S in G'. Since $F^* \le \delta'(S)$ by max-flow min-cut theorem. We have that

$$|A| - c = F^* \le \delta'(S) = \frac{\delta(S)}{\alpha} + |A - S| + \epsilon_\sigma |S - A| = \frac{\delta(S)}{\alpha} + |A| - |A \cap S| + \epsilon_\sigma |S - A|. \tag{3}$$

This implies

$$\frac{\delta(S)}{\alpha} \quad \ge \quad |A \cap S| - \epsilon_\sigma |S - A| - c.$$

Therefore, for any (A,σ)-overlapping cut S where $|S| \ge 10c/\sigma$, we have that

$$\begin{aligned}
\frac{\delta(S)}{|S|} &\ge \alpha\left(\frac{|A \cap S|}{|S|} - \frac{\epsilon_\sigma |S| - |A \cap S|}{|S|} - \frac{c}{|S|}\right) \\
&\ge \alpha((1+\epsilon_\sigma)\sigma - \epsilon_\sigma - \sigma/3) \\
&\ge \alpha(2\sigma/3 - \sigma/3) \\
&\ge \frac{\sigma}{3}\alpha
\end{aligned}$$

where the second inequality is because S is (A,σ)-overlapping and $|S| \ge 10c/\sigma$, and the third inequality is because $\epsilon_\sigma = \frac{1}{3(1/\sigma-1)}$. This concludes the first part of the lemma.

Let S' be the s-t min cut in G'. We have that Equation (3) holds with equality for S'. That is,

$$|A| - c = \frac{\delta(S)}{\alpha} + |A - S| + \epsilon_\sigma |S - A|, \tag{4}$$

[19] In directed graph, an edge e is undirected means that flows can go through e in both direction in contrast to directed edges.

which implies

$$\frac{\delta(S')}{\alpha} = |A \cap S'| - \epsilon_\sigma |S' - A| - c.$$

We want to prove that $\phi(S') = \frac{\delta(S')}{\min\{|S'|,|V-S'|\}} < \alpha$ when $c > 0$. Note again that $\epsilon_\sigma = \frac{1}{3(1/\sigma - 1)} \in [\frac{|A|}{|V-A|}, 1]$ as $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$. First, we have

$$\frac{\delta(S')}{|S'|} = \alpha \frac{|A \cap S'| - \epsilon_\sigma |S' - A| - c}{|S'|} < \alpha.$$

Second, we have

$$\begin{aligned}
\frac{\delta(S')}{|V - S'|} &= \alpha \frac{|A \cap S'| - \epsilon_\sigma |S' - A| - c}{|V - S'|} \\
&= \alpha \frac{(|A| - |A - S'|) - \epsilon_\sigma (|V - A| - |(V - S') - A|) - c}{|V - S'|} \\
&\leq \alpha \frac{\epsilon_\sigma |(V - S') - A| - |A - S'| - c}{|V - S'|} < \alpha,
\end{aligned}$$

where the first inequality is because $\epsilon_\sigma \geq \frac{|A|}{|V-A|}$, and the second inequality is because $\epsilon_\sigma \leq 1$ and $c > 0$. So we conclude that $\phi(S') < \alpha$. Next, we need to show that $|S'|, |V - S'| \geq c$. From Equation (4), we have

$$|A| - c \geq |A - S'| = |A| - |S' \cap A| \implies |S'| \geq |S' \cap A| \geq c.$$

Similarly, we have

$$|A| - c \geq \epsilon_\sigma |S - A| = \epsilon_\sigma |V - A| - \epsilon_\sigma |V - S - A| \geq |A| - \epsilon_\sigma |V - S - A|,$$

which implies

$$|V - S| \geq |V - S - A| \geq c/\epsilon_\sigma \geq c,$$

because $\epsilon_\sigma \leq 1$. This concludes the second part of the lemma. □

In [30], they show how to adjust Goldberg and Rao's binary blocking flow algorithm to *locally* compute s-t min cut in augmented graphs, i.e. the running time does not depend on the size of the whole graph.

Theorem B.25 (Lemma B.3 of [30]). *Given an undirected graph $G = (V, E)$, a set $A \subset V$ where $|A| \leq 4|V - A|$, an expansion parameters α and an overlapping parameter $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$, there is a deterministic algorithm for computing the s-t min cut S of the augmented graph $G_A(\alpha, \sigma)$ in time $O((\frac{vol(A)}{\sigma})^{1.5} \log^2(\frac{vol(A)}{\sigma}))$.*[20]

Given this theorem, we conclude that we our most-balanced sparse cut w.r.t. overlapping cuts algorithm:

[20] In [30], it is shown also that $vol(S) \leq \frac{3}{\sigma} vol(A)$ although we do not need it.

Proof of Theorem 6.5. Given a graph $G = (V, E)$, a set $A \subset V$ where $|A| \leq 4|V - A|$, an expansion parameter α, and an overlapping parameter $\sigma \in [\frac{3|A|}{|V-A|}, \frac{3}{4}]$, we compute the s-t min cut S in the the augmented graph $G_A(\alpha, \sigma)$. Let $c_{size}, c_{exp} = 3/\sigma$.

If the capacity $\delta'(S)$ of S in $G_A(\alpha, \sigma)$ is $|A|$, then by the first part of Lemma B.24, all (A, σ)-overlapping cuts S in G have expansion $\phi(S) \geq \frac{\sigma}{3}\alpha$. Hence, we just report that there is no (α/c_{exp})-sparse (A, σ)-overlapping cuts in G and we are done.

Otherwise, $\delta'(S) = |A| - c$ for some $c > 0$. Then we claim that the s-t min cut S is a (c_{size}, c_{exp})-approximate most-balanced α-sparse cut w.r.t. overlapping cuts where $c_{size}, c_{exp} = 3/\sigma$. Recall $\mathsf{OPT}(G, \alpha/c_{exp}, A, \sigma)$ that is the size of the largest (α/c_{exp})-sparse (A, σ)-overlapping cut S^* where $|S^*| \leq |V - S^*|$. By the first part of Lemma B.24, we have that $\mathsf{OPT}(G, \alpha/c_{exp}, A, \sigma) < 3c/\sigma = c_{size} \cdot c$. But the second part of Lemma B.24 implies that the min-cut S is α-sparse and $|S| \geq c > \mathsf{OPT}(G, \alpha/c_{exp}, A, \sigma)/c_{size}$. Therefore, we return that S is a $(c_{size}, c_{exp}, \alpha, A, \sigma)$-MSO cut of G. □

C Very Simple Dynamic SF with $O(m^{0.5+o(1)})$ Update Time

In this section, let $\beta = 2^{O(\log\log n)^3} = n^{o(1)}$ be the factor from Theorem 4.13. The goal of this section is to show a simple dynamic SF algorithm with update time $O(m^{0.5+o(1)})$ when the size of update sequence is polynomial (the same assumption used in [23]) and the number of edges of the underlying graph is **always bounded** by m:

Theorem C.1. *For any fixed constant c, there is a randomized dynamic SF algorithm for any graphs with n node and at most m edges that works correctly against adaptive adversaries for the first n^c update with high probability and has preprocessing time $O(m\beta \log n)$ and worst-case update time $O(\sqrt{m}\beta \log^2 n)$*[21].

Although there are several faster dynamic SF algorithms, the algorithm from Theorem C.1 is conceptually very simple especially when we use a cut recovery tree from Section 4.2 as a black box.

C.1 The Algorithm

Let $G = (V, E)$ be the underlying graph. To preprocess G, we do the following: 1) construct an arbitrary spanning forest F of G, and 2) preprocess (G, F) using an instance $\mathcal{E}$ of augmented ET tree and an instance $\mathcal{D}$ of cut recovery tree with parameter $\sqrt{m}$.

Remark C.1. In the following description of the algorithm, we only describe how E and F are changed, but we always give the update to $\mathcal{D}$ and $\mathcal{E}$ so that the underlying graphs and forests of them are (G, F).

To insert $e = (u, v)$, we add e into E and if u and v are not connected in F (this can be checked using augmented ET tree), then we also add e into F. To delete $e = (u, v)$, we remove e into E and if $e \in F$ was a tree-edge, then we find a replacement edge e' for reconnecting the two separated components of F as described in Algorithm C.1 and, if e' is returned, add e' into F.

This completes the description of our algorithm.

C.2 Analysis

Let G_i and F_i be the underlying graph and the maintained forest after the i-th update. G_0 and F_0 are the ones after preprocessing. Let $\mathcal{A}$ be our dynamic SF algorithm. Suppose that

[21] To be more precise, the update time is $O((c + c')\sqrt{m}\beta \log^2 n)$ if we want the algorithm to works correctly against adaptive adversaries for the first n^c update with probability $1 - 1/n^{c'}$.

1. Using $\mathcal{E}$, set $T_u = \mathsf{find_tree}(u)$ and $T_v = \mathsf{find_tree}(v)$. If $\mathsf{tree_size}(T_u) < \mathsf{tree_size}(T_v)$, then set $S = V(T_u)$, otherwise $S = V(T_v)$.
2. Sample σ non-tree edges in $G \setminus F$ incident to S using $\mathcal{E}$. If there is a sampled edge e' from $\partial_G(S)$, then RETURN e'.
3. Query k-$\mathsf{cutset}(S)$ to $\mathcal{D}$. If the returned set $\tilde{\partial} \neq \emptyset$, then RETURN a random edge $e' \in \tilde{\partial}$.
4. RETURN "no replacement edge"

Algorithm C.1: The algorithm for finding a replacement edge e' if exists, after deleting an edge $e = (u, v) \in F$.

we fix an adaptive adversary f and a string R which is used as a random choices of $\mathcal{A}$. Let $inv_i(f, \mathcal{A}, R)$ be the event that occurs if F_i is a spanning forest of G_i. For convenience, we write $inv_{<i}(f, \mathcal{A}, R) = inv_0(f, \mathcal{A}, R) \wedge \cdots \wedge inv_{i-1}(f, \mathcal{A}, R)$. Our goal is to prove that the invariant holds with high probability:

Lemma C.2. *Fix any constant c_1 and c_2. Let $\mathcal{A}$ be the dynamic* SF *algorithm from Appendix C and R be the random choices of $\mathcal{A}$. For any adaptive adversary f, we have $\Pr_R[inv_i(f, \mathcal{A}, R)] \geq 1 - 1/n^{c_1}$ for all $0 \leq i \leq n^{c_2}$.*

To prove the above lemma, we need some more lemmas. In Step 2 of Algorithm C.1, suppose that $\delta_G(S) > k$. If no edge $e' \in \partial_{G^\ell}(S^\ell)$ is sampled, then we say that Algorithm C.1 *fails*. Algorithm C.1 is designed such that the following facts holds:

Lemma C.3. *Suppose that Algorithm C.1 does not fail and, before deletion, F was a spanning forest of G. If an edge e' is returned from Algorithm C.1, then e' is a uniformly sampled from the set of replacement edges.*

Proof. Observe that, if F was a spanning forest of G before deletion, then $\partial_G(S)$ is exactly the set of all replacement edges. If $\delta_G(S) > k$, then a random edge e' from $\partial_G(S)$ is returned in Step 2 by the assumption that Algorithm C.1 does not fail. Otherwise, $\delta_G(S) \leq k$. Hence, in Step 3 the queried set returned from k-$\mathsf{cutset}(S)$ is $\tilde{\partial} = \partial_G(S)$ and so a random edge e' from $\partial_G(S)$ is also returned. Therefore, e' is a uniformly sampled from the set of replacement edges. □

Now, the following lemma which immediately implies Lemma C.2.

Lemma C.4. *Let $\mathcal{A}$ be the dynamic* SF *algorithm from Appendix C and R be the random choices of $\mathcal{A}$. For any adaptive adversary f, we have $\Pr_R[inv_i(f, \mathcal{A}, R) \mid inv_{<i}(f, \mathcal{A}, R)] \geq 1 - 1/n^c$ for $i \geq 0$ and c can be made arbitrarily large.*

Proof. When $i = 0$, we have $\Pr_R[inv_i(f, \mathcal{A}, R) \mid inv_{<i}(f, \mathcal{A}, R)] = \Pr_R[inv_0(f, \mathcal{A}, R)] = 1$ because $inv_{<0}(f, \mathcal{A}, R)$ holds trivially and F_0 is a spanning forest of G_0 by construction. When $i \geq 1$, there are two cases for the update at time i: insertion and deletion. First, when we insert $e = (u, v)$, we have $F_i = F_{i-1} \cup \{e\}$ iff u and v are not connected in F_{i-1}. As F_{i-1} is a spanning forest of G_{i-1}, F_i is a spanning forest of G_i. Second, when we delete $e = (u, v)$, suppose that Algorithm C.1 does not *fail*. Since $F_i = F_{i-1} \cup \{e'\} \setminus \{e\}$ and e' is a random replacement edge by Lemma C.3, we again have F_i a spanning forest of G_i as F_{i-1} is a spanning forest of G_{i-1}.

It remains to show that Algorithm C.1 does not fails with high probability. Algorithm C.1 fails when $\delta_G(S) > k$ but no sampled edges in Step 2 are from $\partial_G(S)$. By the property of the

augmented ET-tree $\mathcal{E}$, the probability that a sampled edge is from $\partial_G(S)$ is $\frac{\delta_G(S)}{nt_vol_{(G,F)}(S)}$ for each sampling. Now, the probability that all sampled σ edges are not from $\partial_G(S)$ is at most $(1 - \frac{\delta_G(S)}{nt_vol_{(G,F)}(S)})^\sigma \leq (1 - \frac{\sqrt{m}}{2m})^\sigma \leq 1/n^c$ where c can be made arbitrarily large by choosing the constant in $\sigma = \Theta(\sqrt{m}\log n)$. Note that the first inequality holds because $nt_vol_{(G,F)}(S) \leq 2|E(G)| \leq 2m$ as the number of edges of G is always bounded by m. □

Next, we bound the running time.

Lemma C.5. *Let $\mathcal{A}$ be the dynamic* SF *algorithm from Appendix C. The preprocessing time of $\mathcal{A}$ is $O(m\beta \log n)$, and the update time is $O(\sqrt{m}\beta \log^2 n)$.*

Proof. For the preprocessing time, 1) we can find an arbitrary spanning forest F on an initial graph G in time $O(m)$, and 2) by Theorem 3.3 and Theorem 4.13, we can preprocess (G, F) using augmented ET tree and cut recovery tree in time $O(m\beta \log n)$.

For the update time, we first analyze the running time inside Algorithm C.1. Step 1 takes time at most $O(\log n)$. Step 2 takes time $O(\sigma \log n) = O(\sqrt{m}\log^2 n)$ by Theorem 3.3. Step 3 takes time $O(\sqrt{m}\beta \log^2 n)$ by Theorem 4.13. So the running time in Algorithm C.1 is at most $O(\sqrt{m}\beta \log^2 n)$. The update time spent outside Algorithm C.1 is dominated by the time for updating $\mathcal{D}$. This takes time at most $O(\sqrt{m}\beta \log^2 n)$. □

We conclude that Lemma C.5 and Lemma C.2 together imply Theorem C.1.